# Palo Alto Networks Administrator's Guide

Release 5.0

**Palo Alto Networks, Inc.**
www.paloaltonetworks.com
© 2007-2015 Palo Alto Networks. All rights reserved.
Palo Alto Networks, PAN-OS, and Panorama are trademarks of Palo Alto Networks, Inc. All other trademarks are the property of their respective owners.
P/N 810-000107-00D

# Table of Contents


Preface . . . . . 13

About This Guide . . . . . 13
Organization. . . . . 13
Typographical Conventions. . . . . 15
Notes and Cautions. . . . . 15
Related Documentation . . . . . 15

Chapter 1
Introduction . . . . . 17

Firewall Overview. . . . . 17
Features and Benefits . . . . . 18
Management Interfaces . . . . . 19

Chapter 2
Getting Started . . . . . 21

Preparing the Firewall . . . . . 21
Setting Up the Firewall . . . . . 22
Using the Firewall Web Interface. . . . . 23
Committing Changes . . . . . 25
Navigating to Configuration Pages . . . . . 26
Using Tables on Configuration Pages . . . . . 26
Required Fields . . . . . 26
Locking Transactions . . . . . 27
Supported Browsers . . . . . 27
Getting Help Configuring the Firewall . . . . . 28
Obtaining More Information . . . . . 28
Technical Support . . . . . 28

Chapter 3
Device Management. . . . . 29

System Setup, Configuration, and License Management . . . . . 30
Defining Management Settings . . . . . 30
Defining Operations Settings . . . . . 37
Defining Services Settings . . . . . 41

Defining Content ID Settings . . . . . 43
Defining Session Settings . . . . . 45
SNMP . . . . . 47
Statistics Service . . . . . 48
Comparing Configuration Files . . . . . 49
Installing a License. . . . . 50
Upgrading/Downgrading the PAN-OS Software . . . . . 50
Upgrading PAN-OS in a High Availability Configuration . . . . . 51
Downgrading PAN-OS Software . . . . . 53
Maintenance Release Downgrade . . . . . 53
Feature release Downgrade . . . . . 54
Updating Threat and Application Definitions . . . . . 55
Administrator Roles, Profiles, and Accounts. . . . . 56
Username and Password Requirements. . . . . 57
Defining Administrator Roles . . . . . 58
Defining Password Profiles . . . . . 59
Creating Administrative Accounts . . . . . 59
Specifying Access Domains for Administrators . . . . . 61
Authentication Profiles. . . . . 62
Setting Up Authentication Profiles . . . . . 62
Creating a Local User Database. . . . . 64
Configuring RADIUS Server Settings . . . . . 65
Configuring LDAP Server Settings . . . . . 66
Configuring Kerberos Settings (Native Active Directory Authentication) . . . . 67
Authentication Sequence . . . . . 67
Setting Up Authentication Sequences . . . . . 68
Firewall Logs . . . . . 68
Logging Configuration . . . . . 70
Scheduling Log Exports . . . . . 71
Defining Configuration Log Settings . . . . . 72
Defining System Log Settings . . . . . 72
Defining HIP Match Log Settings . . . . . 73
Defining Alarm Log Settings . . . . . 73
Managing Log Settings . . . . . 74
Configuring SNMP Trap Destinations . . . . . 75
Configuring Syslog Servers. . . . . 76
Custom Syslog Field Descriptions. . . . . 77
Configuring Email Notification Settings. . . . . 83
Viewing Alarms . . . . . 85
Configuring Netflow Settings . . . . . 85
Importing, Exporting and Generating Security Certificates . . . . . 86
Certificates . . . . . 86
Default Trusted Certificate Authorities . . . . . 88
Certificate Profile . . . . . 89
OCSP Responder . . . . . 90
Encrypting Private Keys and Passwords on the Firewall . . . . . 90
Master Key and Diagnostic Settings . . . . . 91
Updating Master Keys. . . . . 91
High Availability . . . . . 93
Active/Passive HA . . . . . 93
Active/Active HA . . . . . 93

Packet Flow . . . . . 94
Deployment Options . . . . . 95
NAT Considerations . . . . . 96
Setting Up HA . . . . . 99
Enabling HA on the Firewall . . . . . 101
Virtual Systems . . . . . 110
Communications Among Virtual Systems . . . . . 111
Shared Gateways . . . . . 112
Defining Virtual Systems . . . . . 113
Configuring Shared Gateways . . . . . 115
Defining Custom Response Pages . . . . . 115
Viewing Support Information . . . . . 117

Chapter 4
Network Configuration . . . . . 119

Firewall Deployment . . . . . 120
Virtual Wire Deployments . . . . . 120
Layer 2 Deployments . . . . . 124
Layer 3 Deployments . . . . . 124
Tap Mode Deployments . . . . . 125
Defining Virtual Wires . . . . . 125
Packet Content Modification . . . . . 126
Firewall Interfaces . . . . . 127
Viewing the Current Interfaces . . . . . 128
Configuring Layer 2 Interfaces . . . . . 128
Configuring Layer 2 Subinterfaces . . . . . 129
Configuring Layer 3 Interfaces . . . . . 130
Configuring Layer 3 Subinterfaces . . . . . 134
Configuring Virtual Wire Interfaces . . . . . 138
Configuring Virtual Wire Subinterfaces . . . . . 139
Configuring Aggregate Interface Groups . . . . . 141
Configuring Aggregate Ethernet Interfaces . . . . . 142
Configuring VLAN Interfaces . . . . . 143
Configuring Loopback Interfaces . . . . . 146
Configuring Tunnel Interfaces . . . . . 147
Configuring Tap Interfaces . . . . . 149
Configuring HA Interfaces . . . . . 150
Security Zones . . . . . 151
Defining Security Zones . . . . . 151
VLAN Support . . . . . 152
Virtual Routers and Routing Protocols . . . . . 153
Routing Information Protocol . . . . . 153
Open Shortest Path First . . . . . 153
Border Gateway Protocol . . . . . 154
Multicast Routing . . . . . 154
Defining Virtual Routers . . . . . 155
DHCP Server and Relay . . . . . 171
DNS Proxy . . . . . 173

Network Profiles . . . . . 174
Defining Interface Management Profiles. . . . . 175
Defining Monitor Profiles. . . . . 177
Defining Zone Protection Profiles . . . . . 178

Chapter 5
Policies and Security Profiles . . . . . 183

Policies. . . . . 183
Guidelines on Defining Policies . . . . . 184
Specifying Users and Applications for Policies . . . . . 186
Security Policies . . . . . 187
Defining Security Policies. . . . . 187
NAT Policies . . . . . 190
Determining Zone Configuration in NAT and Security Policy . . . . . 193
NAT Rule Options . . . . . 193
Defining Network Address Translation Policies . . . . . 194
NAT Policy Examples. . . . . 195
NAT64 . . . . . 196
Policy-Based Forwarding Policies . . . . . 199
Decryption Policies . . . . . 202
Application Override Policies . . . . . 205
Custom Application Definition with Application Override. . . . . 205
Defining Application Override Policies . . . . . 205
Captive Portal Policies . . . . . 206
Defining Captive Portal Policies . . . . . 207
DoS Protection Policies . . . . . 209
Defining DoS Protection Policies . . . . . 209
Security Profiles . . . . . 211
Antivirus Profiles . . . . . 212
Anti-spyware Profiles . . . . . 213
Vulnerability Protection Profiles . . . . . 215
URL Filtering Profiles . . . . . 217
File Blocking Profiles . . . . . 220
Data Filtering Profiles . . . . . 223
DoS Profiles . . . . . 225
Other Policy Objects. . . . . 226
Addresses and Address Groups . . . . . 227
Defining Address Ranges . . . . . 227
Defining Address Groups . . . . . 230
Defining Regions . . . . . 230
Applications and Application Groups . . . . . 231
Defining Applications. . . . . 233
Custom Applications with Signatures . . . . . 236
Defining Application Groups. . . . . 238
Application Filters . . . . . 238
Services . . . . . 239
Service Groups . . . . . 240
Data Patterns . . . . . 240
Custom URL Categories . . . . . 242
Dynamic Block Lists . . . . . 242

Custom Spyware and Vulnerability Signatures . . . . . 243
Defining Data Patterns . . . . . 243
Defining Spyware and Vulnerability Signatures . . . . . 244
Security Profile Groups . . . . . 246
Log Forwarding . . . . . 247
Decryption Profiles . . . . . 248
Schedules . . . . . 250

Chapter 6
Reports and Logs . . . . . 251

Using the Dashboard . . . . . 252
Using the Application Command Center . . . . . 253
Using App-Scope . . . . . 256
Summary Report . . . . . 257
Change Monitor Report . . . . . 258
Threat Monitor Report . . . . . 259
Threat Map Report . . . . . 260
Network Monitor Report . . . . . 261
Traffic Map Report . . . . . 263
Viewing the Logs . . . . . 264
Viewing Session Information . . . . . 267
Working with Botnet Reports . . . . . 267
Configuring the Botnet Report . . . . . 268
Managing Botnet Reports . . . . . 269
Managing PDF Summary Reports . . . . . 270
Managing User Activity Reports . . . . . 272
Managing Report Groups . . . . . 272
Scheduling Reports for Email Delivery . . . . . 273
Viewing Reports . . . . . 273
Generating Custom Reports . . . . . 274
Identifying Unknown Applications and Taking Action . . . . . 276
Taking Action . . . . . 277
Requesting an App-ID from Palo Alto Networks . . . . . 277
Other Unknown Traffic . . . . . 278
Taking Packet Captures . . . . . 278

Chapter 7
Configuring the Firewall for User Identification . . . . . 281

Overview of User Identification . . . . . 281
How User Identification Works . . . . . 281
Identifying Users and Groups . . . . . 282
How User-ID Components Interact . . . . . 283
User-ID Agent . . . . . 283
PAN-OS User Mapping . . . . . 283
Terminal Services Agent . . . . . 283
PAN-OS LDAP Group Query . . . . . 284

User Identification Agents . . . . . 284
Captive Portals . . . . . 285
Configuring the Firewall for User Identification . . . . . 286
PAN-OS User Mapping Configuration . . . . . 291
Configuring PAN-OS User Mapping . . . . . 292
Configure a Firewall to Share User Mapping Data . . . . . 295
Setting Up the User-ID Agent . . . . . 296
Installing the User-ID Agent . . . . . 297
Configuring the User-ID Agent . . . . . 298
Discovering Domain Controllers . . . . . 301
Monitoring User-ID Agent Operation . . . . . 301
Uninstalling and Upgrading the User-ID Agent . . . . . 301
Setting Up the Terminal Services Agent . . . . . 301
Installing or Upgrading the Terminal Server Agent on the Terminal Server . 302
Configuring the Terminal Server Agent on the Terminal Server . . . . . 303
Uninstalling the Terminal Server Agent on the Terminal Server . . . . . 307

Chapter 8
Configuring IPSec Tunnels . . . . . 309

Virtual Private Networks . . . . . 310
VPN Tunnels . . . . . 311
IPSec and IKE . . . . . 311
IPSec and IKE Crypto Profiles . . . . . 312
Setting Up IPSec VPNs . . . . . 313
Defining IKE Gateways . . . . . 314
Setting Up IPSec Tunnels . . . . . 315
Defining IKE Crypto Profiles . . . . . 318
Defining IPSec Crypto Profiles . . . . . 319
Viewing IPSec Tunnel Status on the Firewall . . . . . 319
Sample VPN Configuration . . . . . 320
Existing Topology . . . . . 320
New Topology . . . . . 321
Configure the VPN Connection . . . . . 321
VPN Connectivity Troubleshooting . . . . . 322
GlobalProtect Large Scale VPN Deployment . . . . . 323
Overview . . . . . 323
Deploying a Large Scale VPN Network . . . . . 324
Certificates and the OCSP Responder . . . . . 325
Global Protect Gateway Configuration . . . . . 328
GlobalProtect Portal Configuration . . . . . 330
GlobalProtect Satellite Configuration . . . . . 332
Dynamic Routing Protocols and Large Scale VPNs . . . . . 333
Backing up a GlobalProtect Portal . . . . . 334

Chapter 9
Configuring GlobalProtect . . . . . 335
Overview . . . . . 335
GlobalProtect Authentication . . . . . 336
Setting Up GlobalProtect . . . . . 337
Setting Up and Activating the GlobalProtect Agent . . . . . 352
Setting Up the GlobalProtect Agent . . . . . 353

Chapter 10
Configuring Quality of Service. . . . . 355
Firewall Support for QoS . . . . . 355
Configuring QoS for Firewall Interfaces . . . . . 356
Defining QoS Profiles . . . . . 358
Defining QoS Policies . . . . . 359
Displaying QoS Statistics . . . . . 362

Chapter 11
Setting Up a VM-Series Firewall . . . . . 363
Overview . . . . . 363
System Requirements and Limitations . . . . . 364
Requirements . . . . . 364
Limitations . . . . . 364
About Licensing the VM-Series Firewall . . . . . 365
Registering the VM-Series Firewall. . . . . 365
Installing and Licensing the VM-Series Firewall . . . . . 366
Troubleshooting . . . . . 370

Chapter 12
Setting Up Panorama . . . . . 371
Overview . . . . . 371
Setting Up Panorama as a Virtual Appliance . . . . . 372
Installing Panorama . . . . . 372
Configuring the Panorama Network Interface . . . . . 373
Expanding the Log Storage Capacity . . . . . 373
Adding a Virtual Disk . . . . . 374
Setting Up Storage Partitions . . . . . 374
Setting up Panorama on an M-Series Appliance . . . . . 376
Performing Initial Setup . . . . . 376
Logging in to Panorama . . . . . 377
Changing the Default Password . . . . . 377
Configuring High Availability (HA) . . . . . 377
Switching the Logging Priority in an HA Pair . . . . . 379

Chapter 13
Central Device Management Using Panorama . . . . . 381

Accessing the Panorama Web Interface . . . . . 382
Using the Panorama Interface . . . . . 382
Panorama Tab . . . . . 383
Adding Devices . . . . . 385
Defining Device Groups . . . . . 386
Panorama Administrator Roles, Profiles, and Accounts . . . . . 387
Defining Panorama Administrator Roles . . . . . 388
Creating Panorama Administrative Accounts . . . . . 389
Specifying Panorama Access Domains for Administrators . . . . . 392
Device Groups . . . . . 393
Working with Policies . . . . . 393
Working with Objects . . . . . 395
Working with Devices . . . . . 397
Commit Operation in Panorama . . . . . 398
Panorama Backward Compatibility . . . . . 399
Templates . . . . . 399
Configuring Panorama Templates . . . . . 401
Adding a New Template . . . . . 401
Configuring a Template . . . . . 402
Overriding Template Settings . . . . . 402
Removing Templates . . . . . 403
Logging . . . . . 403
Logging and Reporting . . . . . 403
Generating User Activity Reports . . . . . 404
Using Panorama for Log Collection . . . . . 404
Deploying Distributed Log Collection . . . . . 405
Managing Log Collectors . . . . . 407
Defining Log Collector Groups . . . . . 411
Viewing Firewall Deployment Information . . . . . 416
Backing Up Firewall Configurations . . . . . 417
Scheduling Configuration Exports . . . . . 417
Upgrading the Panorama Software . . . . . 419

Chapter 14
Configuring WildFire . . . . . 421

About WildFire . . . . . 421
Setting Up WildFire on the Firewall . . . . . 423
Configuring WildFire Settings on the Firewall . . . . . 423
Configuring WildFire Forwarding . . . . . 424
WildFire Data Filtering Log . . . . . 425
Using the WildFire Portal . . . . . 425
Configuring Settings on the WildFire Portal . . . . . 426
Viewing WildFire Reports . . . . . 427

Appendix A
Custom Pages . . . . . 429

Default Antivirus Response Page . . . . . 429
Default Application Block Page . . . . . 431
Default File Blocking Block Page . . . . . 431
Default URL Filtering Response Page . . . . . 432
Default Anti-spyware Download Response Page . . . . . 433
Default Decryption Opt-out Response Page . . . . . 433
Captive Portal Comfort Page . . . . . 434
URL Filtering Continue and Override Page . . . . . 434
SSL VPN Login Page . . . . . 435
SSL Certificate Revoked Notify Page . . . . . 436

Appendix B
Application Categories, Subcategories, Technologies, and Characteristics 437

Application Categories and Subcategories . . . . . 437
Application Technologies . . . . . 439
Application Characteristics . . . . . 439

Appendix C
Common Criteria/Federal Information Processing Standards Support . 441

Enabling CC/FIPS Mode . . . . . 441
CC/FIPS Security Functions . . . . . 443

Appendix D
Open Source Licenses . . . . . 445

Artistic License . . . . . 446
BSD . . . . . 447
GNU General Public License. . . . . 448
GNU Lesser General Public License . . . . . 452
MIT/X11 . . . . . 458
OpenSSH . . . . . 458
PSF . . . . . 462
PHP . . . . . 462
Zlib . . . . . 463

Index. . . . . 465

# Preface

This preface contains the following sections:


- "About This Guide" in the next section
- "Organization" on page 13
- "Typographical Conventions" on page 15
- "Notes and Cautions" on page 15
- "Related Documentation" on page 15


## About This Guide

This guide describes how to administer the Palo Alto Networks firewall using the device's web interface.

This guide is intended for system administrators responsible for deploying, operating, and maintaining the firewall.

## Organization

This guide is organized as follows:

- **Chapter 1, "Introduction"**—Provides an overview of the firewall.
- **Chapter 2, "Getting Started"**—Describes how to install the firewall.
- **Chapter 3, "Device Management"—**Describes how to perform basic system configuration and maintenance for the firewall, including how to configure a pair of firewalls for high availability, define user accounts, update the software, and manage configurations.
- **Chapter 4, "Network Configuration"**—Describes how to configure the firewall for your network, including routing configuration.
- **Chapter 5, "Policies and Security Profiles"**—Describes how to configure security policies and profiles by zone, users, source/destination address, and application.
- **Chapter 6, "Reports and Logs"**—Describes how to view the reports and logs provided with the firewall.

- **Chapter 7, "Configuring the Firewall for User Identification"**—Describes how to configure the firewall to identify the users who attempt to access the network.
- **Chapter 8, "Configuring IPSec Tunnels"**—Describes how to configure IP Security (IPSec) tunnels on the firewall.
- **Chapter 9, "Configuring GlobalProtect"**—Describes GlobalProtect, which allows secure login from client systems located anywhere in the world.
- **Chapter 10, "Configuring Quality of Service"**—Describes how to configure quality of service (QoS) on the firewall.
- **Chapter 12, "Setting Up Panorama"**—Describes how to install the centralized management system for the Palo Alto Networks firewall.
- **Chapter 13, "Central Device Management Using Panorama"**—Describes how to use Panorama to manage multiple firewalls.
- **Chapter 14, "Configuring WildFire"**—describes how to use WildFire for analysis and reporting on malware that traverses the firewall.
- **Appendix A, "Custom Pages"**—Provides HTML code for custom response pages to notify end users of policy violations or special access conditions.
- **Appendix B, "Application Categories, Subcategories, Technologies, and Characteristics"**—Contains a list of the application categories defined by Palo Alto Networks.
- **Appendix C, "Common Criteria/Federal Information Processing Standards Support"**—Describes firewall support for the Federal Information Processing Standards 140-2.
- **Appendix D, "Open Source Licenses"**—Includes information on applicable open source licenses.

# Typographical Conventions

This guide uses the following typographical conventions for special terms and instructions.

| Convention | Meaning | Example |
|---|---|---|
| **boldface** | Names of commands, keywords, and selectable items in the web interface | Click **Security** to open the Security Rules page. |
| *italics* | Name of parameters, files, directories, or Uniform Resource Locators (URLs) | The address of the Palo Alto Networks home page is *http://www.paloaltonetworks.com* |
| `courier font` | Coding examples and text that you enter at the command prompt | Enter the following command:<br>`set deviceconfig system dns-settings` |
| Click | Click the left mouse button | Click **Administrators** under the **Devices** tab. |
| Right-click | Click the right mouse button. | Right-click on the number of a rule you want to copy, and select **Clone Rule**. |

# Notes and Cautions

This guide uses the following symbols for notes and cautions.

| Symbol | Description |
|---|---|
| | NOTE<br>Indicates helpful suggestions or supplementary information. |
| | CAUTION<br>Indicates actions that could cause loss of data. |

# Related Documentation

The following additional documentation is provided with the firewall:

- *Getting Started Guide*
- *Quick Start*
- *Palo Alto Networks License Agreement and Warranty*

You can find additional related documentation at https://live.paloaltonetworks.com/community/documentation.

# Chapter 1
# Introduction

This chapter provides an overview of the firewall:


- "Firewall Overview" in the next section
- "Features and Benefits" on page 18
- "Management Interfaces" on page 19


## Firewall Overview

The Palo Alto Networks firewall allows you to specify security policies based on accurate identification of each application seeking access to your network. Unlike traditional firewalls that identify applications only by protocol and port number, the firewall uses packet inspection and a library of application signatures to distinguish between applications that have the same protocol and port, and to identify potentially malicious applications that use non-standard ports.

For example, you can define security policies for specific applications, rather than rely on a single policy for all port 80 connections. For each identified application, you can specify a security policy to block or allow traffic based on the source and destination zones and addresses (IPv4 and IPv6). Each security policy can also specify security profiles to protect against viruses, spyware, and other threats.

# Features and Benefits

The firewall provides granular control over the traffic allowed to access your network. The primary features and benefits include:

- **Application-based policy enforcement**—Access control by application is far more effective when application identification is based on more than just protocol and port number. High risk applications can be blocked, as well as high risk behavior, such as file-sharing. Traffic encrypted with the s Layer (SSL) protocol can be decrypted and inspected.
- **User Identification (User-ID)**—User-ID allows administrators to configure and enforce firewall policies based on users and user groups, instead of or in addition to network zones and addresses. The firewall can communicate with many directory servers, such as Microsoft Active Directory, eDirectory, SunOne, OpenLDAP, and most other LDAP based directory servers to provide user and group information to the firewall. This information can then be used to provide an invaluable method of providing secure application enablement that can be defined per user or group. For example, the administrator could allow one organization to use a web-based application, but no other organizations in the company would be able to use that application. You can also configure granular control of certain components of an application based on users and groups. Refer to "Configuring the Firewall for User Identification" on page 281.
- **Threat prevention**—Threat prevention services that protect the network from viruses, worms, spyware, and other malicious traffic can be varied by application and traffic source (refer to "Security Profiles" on page 211).
- **URL filtering**—Outbound connections can be filtered to prevent access to inappropriate web sites (refer to "URL Filtering Profiles" on page 217).
- **Traffic visibility**—Extensive reports, logs, and notification mechanisms provide detailed visibility into network application traffic and security events. The Application Command Center (ACC) in the web interface identifies the applications with the most traffic and the highest security risk (refer to "Reports and Logs" on page 251).
- **Networking versatility and speed**—The firewall can augment or replace your existing firewall, and can be installed transparently in any network or configured to support a switched or routed environment. Multi-gigabit speeds and a single-pass architecture provide all services with little or no impact on network latency.
- **GlobalProtect**—GlobalProtect provides security for client systems, such as laptops, that are used in the field by allowing easy and secure login from anywhere in the world.
- **Fail-safe operation**—High availability support provides automatic failover in the event of any hardware or software disruption (refer to "Enabling HA on the Firewall" on page 101).
- **Malware analysis and reporting**—WildFire provides detailed analysis and reporting on malware that traverses the firewall.
- **VM-Series Firewall**—Provides a virtual instance of PAN-OS positioned for use in a virtualized data center environment and particularly well suited for private and public cloud deployments. Installs on any x86 device that is capable of running VMware ESXi, without the need to deploy Palo Alto Networks hardware.

- **Management and Panorama—**Each firewall is managed through an intuitive web interface or a command-line interface (CLI), or all devices can be centrally managed through the Panorama centralized management system, which has a web interface very similar to the device web interface.

# Management Interfaces

The firewall supports the following management interfaces. Refer to "Supported Browsers" on page 27 for a list of supported browsers.

- **Web interface**—Configuration and monitoring over HTTP or HTTPS from a web browser.
- **CLI**—Text-based configuration and monitoring over Telnet, Secure Shell (SSH), or the console port (refer to the *PAN-OS Command Line Interface Reference Guide*).
- **Panorama**—Palo Alto Networks product that provides web-based management, reporting, and logging for multiple firewalls. The Panorama interface is similar to the device web interface, with additional management functions included. Refer to "Setting Up Panorama" on page 371 for instructions on installing Panorama and "Central Device Management Using Panorama" on page 381 for information on using Panorama.
- **Simple Network Management Protocol (SNMP)**—Palo Alto Networks products support SNMPv2c and SNMPv3, read-only access over SNMP, and support for TRAPS. Refer to "Configuring SNMP Trap Destinations" on page 75).
- **Syslog**—Provides message generation for one or more remote syslog servers (refer to "Configuring Syslog Servers" on page 76).
- **XML API**—Provides a Representational State Transfer (REST)-based interface to access device configuration, operational status, reports, and packet captures from the firewall. There is an API browser available on the firewall at *https://<firewall>/api*, where *<firewall>* is the host name or IP address of the firewall. This link provides help on the parameters required for each type of API call. An XML API usage guide is available on the DevCenter online community at *http://live.paloaltonetworks.com*.

# Chapter 2
# Getting Started

This chapter describes how to set up and start using the firewall:

- "Preparing the Firewall" in the next section
- "Setting Up the Firewall" on page 22
- "Using the Firewall Web Interface" on page 23
- "Getting Help Configuring the Firewall" on page 28

***Note:*** *Refer to "Setting Up Panorama" on page 371 for instructions on installing the Panorama centralized management system.*

## Preparing the Firewall

Perform the following tasks to prepare the firewall for setup:

1. Mount the firewall in a rack and power it up as described in the *Hardware Reference Guide*.
2. Register your firewall at *https://support.paloaltonetworks.com* to obtain the latest software and App-ID updates, and to activate support or subscriptions with the authorization codes emailed to you.
3. Obtain an IP address from your network administrator for configuring the management port on the firewall.

# Setting Up the Firewall

To perform the initial firewall setup:

1. Connect your computer to the management port (MGT) on the firewall using an RJ-45 Ethernet cable.
2. Start your computer. Assign a static IP address to your computer on the 192.168.1.0 network (for example, 192.168.1.5) with a netmask of 255.255.255.0.
3. Launch a supported web browser and enter **https://192.168.1.1**.

   The browser automatically opens the Palo Alto Networks login page.
4. Enter **admin** in both the **Name** and **Password** fields, and click **Login**. The system presents a warning that the default password should be changed. Click **OK** to continue.
5. On the **Device** tab, choose **Setup** and configure the following (for general instructions on configuring settings in the web interface, refer to "Using the Firewall Web Interface" on page 23):
   - On the **Management** tab under **Management Interface Settings**, enter the firewall's IP address, netmask, and default gateway.
   - On the **Services** tab, enter the IP address of the Domain Name System (DNS) server. Enter the IP address or host and domain name of the Network Time Protocol (NTP) server and select your time zone.
   - Click **Support** on the side menu.
     If this is the first Palo Alto Networks firewall for your company, click **Register Device** to register the firewall. (If you have already registered a firewall, you have received a user name and password.)
     Click the **Activate support using authorization codes** link and enter the authorization codes that have been emailed to you for any optional features. Use a space to separate multiple authorization codes.
6. Click **Administrators** under the **Devices** tab.
7. Click **admin**.
8. In the **New Password** and **Confirm New Password** fields, enter and confirm a case-sensitive password (up to 15 characters).
9. Click **OK** to submit the new password.
10. Commit the configuration to make these settings active. When the changes are committed, the firewall will be reachable through the IP address assigned in Step 5. For information on committing changes, refer to "Committing Changes" on page 25.

***Note:*** *The default configuration of the firewall when delivered from the factory, or after a factory reset is performed, is a virtual wire between Ethernet ports 1 and 2 with a default policy to deny all inbound traffic and allow all outbound traffic.*

# Using the Firewall Web Interface

The following conventions apply when using the firewall interface.

- To display the menu items for a general functional category, click the tab, such as **Objects** or **Device**, near the top of the browser window.

- Click an item on the side menu to display a panel.

- To display submenu items, click the icon to the left of an item. To hide submenu items, click the icon to the left of the item.

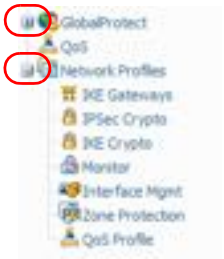


- On most configuration pages, you can click **Add** to create a new item.

- To delete one or more items, select their check boxes and click **Delete**. In most cases, the system prompts you to confirm by clicking **OK** or to cancel the deletion by clicking **Cancel**.

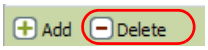


- On some configuration pages, you can select the check box for an item and click **Clone** to create a new item with the same information as the selected item.

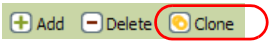

- To modify an item, click its underlined link.

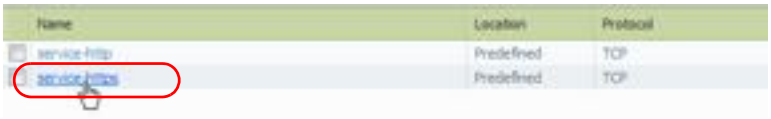


- To view help information on a page, click the **Help** icon in upper right area of the page.

- To view the current list of tasks, click the **Tasks** icon in the lower right corner of the page. The Task Manager window opens to show the list of tasks, along with status, start times, associated messages, and actions. Use the **Show** drop-down list to filter the list of tasks.

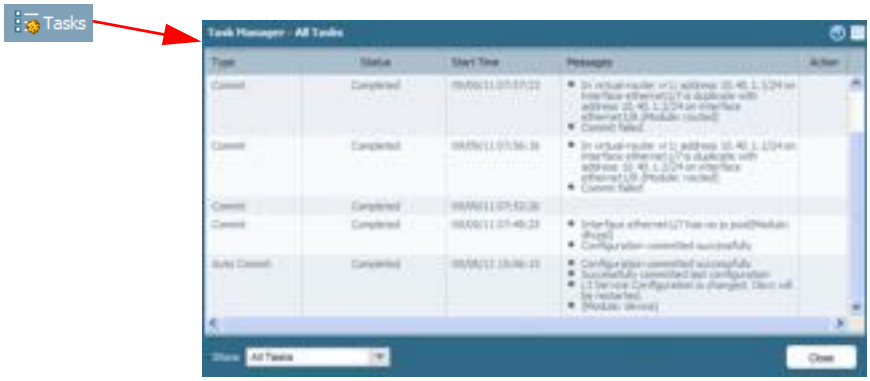


- The web interface language is controlled by the current language of the computer that is managing the device if a specific language preference has not been defined. For example, if the computer you use to manage the firewall has a locale of Spanish, when you log in to the firewall, the web interface will be in Spanish.

  To specify a language that will always be used for a given account regardless of the locale of the computer, click the **Language** icon in the lower right corner of the page and the Language Preference window opens. Click the drop-down list to select the desired language and then click **OK** to save your change.

- On pages that list information you can modify (for example, the **Setup** page on the **Devices** tab), click the icon in the upper right corner of a section to edit the settings.

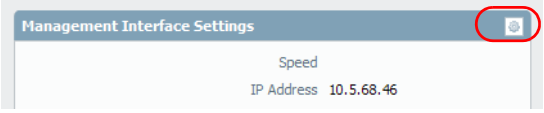


- After you configure settings, you must click **OK** or **Save** to store the changes. When you click **OK**, the current "candidate" configuration is updated.

## Committing Changes

Click **Commit** at the top of the web interface to open the commit dialog box.

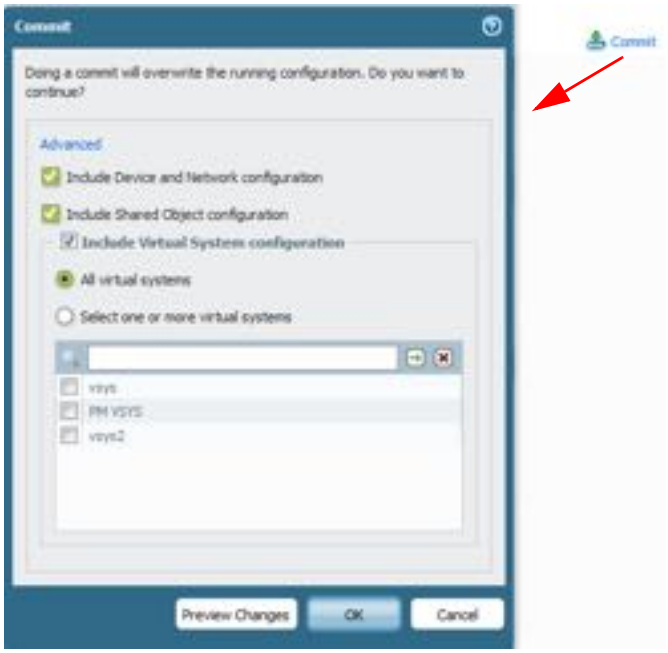


The following options are available in the commit dialog box. Click the **Advanced** link, if needed, to display the options:

- **Include Device and Network configuration**—Include the device and network configuration changes in the commit operation.
- **Include Shared Object configuration**—(Multi-virtual system firewalls only) Include the shared object configuration changes in the commit operation.
- **Include Policy and Objects**—(Non-multi-virtual system firewalls only) Include the policy and object configuration changes in the commit operation.
- **Include virtual system configuration**—Include all virtual systems or choose **Select one or more virtual systems**.

  For more information about committing changes, refer to "Defining Operations Settings" on page 37.
- **Preview Changes**—Click this button to bring up a two-pane window that shows proposed changes in the candidate configuration compared to the current running configuration. You can choose the number of lines of context to display, or show all lines. Changes are color coded based on items that have been added, modified, or

deleted.
The **Device > Config Audit** feature performs the same function, refer to "Comparing Configuration Files" on page 49.

***Note:*** *Configuration changes that span multiple configuration areas may require a full commit. For example, if you click* ***Commit*** *and only select the* ***Include Device and Network configuration*** *option, some items that you changed in the Device tab will not commit. This includes certificates and User-ID options as well as Server Profiles used for User-ID, such as an LDAP server profile. This can also occur if you perform a partial commit after importing a configuration. To commit these types of changes, do a full commit and select both* ***Include Device and Network configuration and Include Policy*** *and* ***Object configuration****.*

## Navigating to Configuration Pages

Each configuration section in this guide shows the menu path to the configuration page. For example, to reach the **Vulnerability Protection** page, choose the **Objects** tab and then choose **Vulnerability Protection** under **Security Profiles** in the side menu. This is indicated in this guide by the following path:

▶ ***Objects > Security Profiles > Vulnerability Protection***

## Using Tables on Configuration Pages

The tables on configuration pages include sorting and column chooser options. Click a column header to sort on that column, and click again to change the sort order. Click the arrow to the right of any column and select check boxes to choose the columns to display.

## Required Fields

Required fields are shown with a light yellow background. A message indicating that the field is required appears when you hover over or click in the field entry area.

# Locking Transactions

The web interface provides support for multiple administrators by allowing an administrator to lock a current set of transactions, thereby preventing configuration changes or commit operations by another administrator until the lock is removed. The following types of locks are supported:

- **Config lock**—Blocks other administrators from making changes to the configuration. This type of lock can be set globally or for a virtual system. It can be removed only by the administrator who set it or by a superuser on the system.
- **Commit Lock**—Blocks other administrators from committing changes until all of the locks have been released. This type of lock prevents collisions that can occur when two administrators are making changes at the same time and the first administrator finishes and commits changes before the second administrator has finished. The lock is released when the current changes are committed by the administrator who applied the lock, or it can be released manually.

Any administrator can open the lock window to view the current transactions that are locked, along with a timestamp for each.

To lock a transaction, click the unlocked icon on the top bar to open the Locks dialog box. Click **Take a Lock**, select the scope of the lock from the drop-down list, and click **OK**. Add additional locks as needed, and then click **Close** to close the Lock dialog box.

The transaction is locked, and the icon on the top bar changes to a locked icon that shows the number of locked items in parentheses.

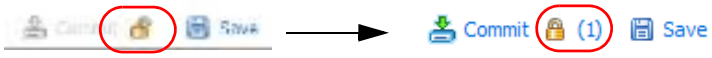


To unlock a transaction, click the locked icon on the top bar to open the Locks window. Click the icon for the lock that you want to remove, and click **Yes** to confirm. Click **Close** to close the Lock dialog box.

You can arrange to automatically acquire a commit lock by selecting the **Automatically acquire commit lock** check box in the Management area of the **Device Setup** page. Refer to "System Setup, Configuration, and License Management" on page 30.

# Supported Browsers

The following web browsers are supported for access to the firewall web interface:

- Internet Explorer 7+
- Firefox 3.6+
- Safari 5+
- Chrome 11+

# Getting Help Configuring the Firewall

Use the information in this section to obtain help on using the firewall.

## Obtaining More Information

To obtain more information about the firewall, refer to the following:

- **General information**—Go to *http://www.paloaltonetworks.com*.
- **Online help**—Click **Help** in the upper-right corner of the web interface to access the online help system.
- **Collaborative area for customer/partner interaction to share tips, scripts, and signatures**—Go to *https://live.paloaltonetworks.com/community/devcenter.*

## Technical Support

For technical support, use the following methods:

- Go to the KnowledgePoint online support community at *http://live.paloaltonetworks.com.*
- Go to *https://support.paloaltonetworks.com*.

## Chapter 3
# Device Management

This chapter describes how to perform basic system configuration and maintenance for the firewall and includes overviews of the virtual systems, high availability, and logging functions:


- "System Setup, Configuration, and License Management" in the next section
- "Comparing Configuration Files" on page 49
- "Installing a License" on page 50
- "Upgrading/Downgrading the PAN-OS Software" on page 50
- "Updating Threat and Application Definitions" on page 55
- "Administrator Roles, Profiles, and Accounts" on page 56
- "Authentication Profiles" on page 62
- "Authentication Sequence" on page 67
- "Certificate Profile" on page 89
- "Firewall Logs" on page 68
- "Configuring SNMP Trap Destinations" on page 75
- "Configuring Syslog Servers" on page 76
- "Configuring Email Notification Settings" on page 83
- "Viewing Alarms" on page 85
- "Configuring Netflow Settings" on page 85
- "Importing, Exporting and Generating Security Certificates" on page 86
- "High Availability" on page 93
- "Virtual Systems" on page 110
- "Defining Custom Response Pages" on page 115
- "Viewing Support Information" on page 117

# System Setup, Configuration, and License Management

The following sections describe how to define the network settings and manage configurations for the firewall:


- "Defining Management Settings" in the next section
- "Defining Operations Settings" on page 37
- "Defining Services Settings" on page 41
- "Defining Content ID Settings" on page 43
- "Defining Session Settings" on page 45
- "SNMP" on page 47
- "Statistics Service" on page 48
- "Installing a License" on page 50

***Note:*** *Refer to "Configuring WildFire" on page 421 for information on configuring the settings on the WildFire tab.*

## Defining Management Settings

▶ ***Device > Setup > Management***

The **Setup** page allows you to configure the firewall for management, operations, services, content identification, WildFire malware analysis and reporting, and session behavior.

If you do not want to use the management port, you can define a loopback interface and manage the firewall through the IP address of the loopback interface (refer to "Configuring Loopback Interfaces" on page 146).

Perform any of the following operations on this page:

- To change the host name or network settings, click **Edit** on the first table on the page, and specify the following information.

**Table 1. Management Settings**

| Item | Description |
|---|---|
| **General Settings** | |
| Hostname | Enter a host name (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Domain | Enter the Fully Qualified Domain Name (FQDN) of the firewall (up to 31 characters). |
| Login Banner | Enter custom text that will be displayed on the firewall login page. The text is displayed below the **Name** and **Password** fields. |

**Table 1. Management Settings (Continued)**

| Item | Description |
|---|---|
| Time Zone | Select the time zone of the firewall. |
| Locale | Select a language for PDF reports from the drop-down list. Refer to "Managing PDF Summary Reports" on page 270.<br>If you have a specific language preference set for the web interface, PDF reports will still use the language specified in this locale setting. Refer to language preference in "Using the Firewall Web Interface" on page 23. |
| Time | To set the date and time on the firewall, click **Set Time**. Enter the current date in (YYYY/MM/DD) or click the calendar icon to select a month and day. Enter the current time in 24-hour format (HH:MM:SS). You can also define an NTP server from Device > Setup > Services. |
| Serial Number | (Panorama only) Enter the serial number of the firewall. |
| Geo Location | Enter the latitude (-90.0 to 90.0) and longitude (-180.0 to 180.0) of the firewall. |
| Automatically acquire commit lock | Automatically apply a commit lock when you change the candidate configuration. For more information, refer to "Locking Transactions" on page 27. |
| Certificate Expiration Check | Instruct the firewall to create warning messages when on-box certificates near their expiration dates. |
| Multi Virtual System Capability | To enable the use of multiple virtual systems (if supported on the firewall model), click **Edit** for Multi Virtual System Capability near the top of the **Setup** page. Select the check box, and click **OK**. For more information about virtual systems, refer to "Virtual Systems" on page 110. |
| **Authentication Settings** | |
| Authentication Profile | Select the authentication profile to use for administrator access to the firewall. For instructions on configuring authentication profiles, refer to "Setting Up Authentication Profiles" on page 62. |
| Certificate Profile | Select the certificate profile to use for administrator access to the firewall. For instructions on configuring certificate profiles, refer to "Certificate Profile" on page 89. |
| Idle Timeout | Enter the timeout interval (1 - 1440 minutes). A value of 0 means that the management, web, or CLI session does not time out. |
| # Failed Attempts | Enter the number of failed login attempts that are allowed for the web interface and CLI before the account is locked. (1-10, default 0). 0 means that there is no limit. |
| Lockout Time | Enter the number of minutes that a user is locked out (0-60 minutes) if the number of failed attempts is reached. The default 0 means that there is no limit to the number of attempts. |

**Table 1. Management Settings (Continued)**

| Item | Description |
|---|---|
| **Panorama Settings** | |
| Panorama Servers | Enter the IP address of Panorama, the Palo Alto Networks centralized management system (if any). The server address is required to manage the device using Panorama. If Panorama is in an HA configuration, enter the secondary Panorama server IP address in the second Panorama Servers field.<br>***Note:*** *To remove any policies that Panorama propagates to managed firewalls, click the* ***Disable Panorama Policy and Objects*** *link. To keep a local copy of the policies and objects to your device before removing them from Panorama, click the* ***Import Panorama Policy and Objects before disabling*** *check box in the dialog box that opens. Click* ***OK****.*<br>***Note:*** *When you select the import check box, the policies and objects will be copied to the current candidate configuration. If you commit this configuration, the policies and objects will become part of your configuration and will no longer be managed by Panorama.*<br>To remove device and network templates, click the **Disable Device and Network Template** link. To keep a local copy of the device and network templates, click the **Import Device and Network Templates before disabling** check box in the dialog box that opens and click **OK**. When you select the import check box, the configuration defined in the device and network templates will be copied to the current candidate configuration. If you commit that configuration, these items will become part of your configuration and will no longer be managed by Panorama. Templates will no longer be accepted on the device until you click **Enable Device and Network Templates**. |
| Receive Timeout for connection to device/ Panorama | • Device—Enter the timeout for receiving TCP messages from all managed devices (1-240 seconds, default 240).<br>• Panorama—Enter the timeout for receiving TCP messages from Panorama (1-240 seconds, default 240). |
| Send Timeout for connection to device/ Panorama | • Device—Enter the timeout for receiving TCP messages from all managed devices (1-240 seconds, default 240).<br>• Panorama—Enter the timeout for receiving TCP messages from Panorama (1-240 seconds, default 240). |
| Retry Count for SSL send to device/ Panorama | • Device—Enter the number of retries for attempts to send Secure Socket Layer (SSL) messages to managed devices (1-64, default 25).<br>• Panorama—Enter the number of retries for attempts to send Secure Socket Layer (SSL) messages to Panorama (1-64, default 25). |
| Share Unused Address and Service Objects with Devices (Panorama only) | Select this check box to share all Panorama shared objects and device group specific objects with managed devices. When unchecked, Panorama policies are checked for references to address, address group, service, and service group objects and any objects that are not referenced will not be shared. This option will ensure that only necessary objects are being sent to managed devices in order to reduce the total object count. |

**Table 1. Management Settings (Continued)**

| Item | Description |
|---|---|
| Shared Objects Take Precedence (Panorama only) | Select the check box to specify that shared objects take precedence over device group objects. This option is a system-wide setting and is off by default. When this option is off, device groups override corresponding objects of the same name. If the option is selected, device group objects cannot override corresponding objects of the same name from a shared location and any device group object with the same name as a shared object will be discarded. |
| **Management Interface Settings** | |
| MGT Interface Speed | Configure a data rate and duplex option for the management interface. The choices include 10Mbps, 100Mbps, and 1Gbps at full or half duplex. Use the default auto-negotiate setting to have the firewall determine the interface speed.<br>This setting should match the port settings on the neighboring network equipment. |
| MGT Interface IP Address | Enter the IP address of the management port. Alternatively, you can use the IP address of a loopback interface for device management. This address is used as the source address for remote logging. |
| Netmask | Enter the network mask for the IP address, such as "255.255.255.0". |
| Default Gateway | Enter the IP address of the default router (must be on the same subnet as the management port). |
| MGT Interface IPv6 Address | (Optional) Enter the IPv6 address of the management port. An IPv6 prefix length is required to indicate the netmask, for example 2001:400:f00::1/64. |
| Default IPv6 Gateway | Enter the IPv6 address of the default router (must be on the same subnet as the management port), if you assigned an IPv6 address to the management port. |
| MGT Interface Services | Select the services enabled on the specified management interface address: HTTP, HTTPS, Telnet, Secure Shell (SSH), and/or ping. |
| Permitted IPs | Enter the list of IP addresses from which firewall management is allowed. When using this option for Panorama, you will need to make sure that each managed device has it's IP address added, otherwise it will not be able to connect send logs to Panorama or receive configuration updates. |

**Table 1. Management Settings (Continued)**

| Item | Description |
|---|---|
| **Logging and Reporting Settings** | |
| Log Storage | Specify the percentage of space allocated to each log type on the hard disk.<br>When you change a percent value, the associated disk allocation changes automatically. If the total of all the values exceeds 100%, a message appears on the page in red, and an error message is presented when you attempt to save the settings. If this occurs, readjust the percentages so the total is within the 100% limit.<br>Click **OK** to save settings and **Restore Defaults** to restore all of the default settings.<br>***Note:*** *When a log reaches its maximum size, it starts to be overwritten beginning with the oldest entries. If you resize an existing log to be smaller than its current size, the firewall starts immediately to cut down the log when you commit the changes, with the oldest logs removed first.* |
| Max Rows in User Activity Report | Enter the maximum number of rows that is supported for the detailed user activity reports (1-1048576, default 65535). |
| Max Rows in CSV Export | Enter the maximum number of rows that will appear in the CSV reports generated from the **Export to CSV** icon in the traffic logs view (range 1-1048576, default 65535). |
| Number of Versions for Config Audit | Enter the number of configuration audit versions to save before discarding the oldest ones (default 100). |
| Number of Versions for Config Backups | (Panorama only) Enter the number of configuration backups to save before discarding the oldest ones (default 100). |
| Average Browse Time (sec) | Configure this variable to adjust how browse time is calculated in the "User Activity Report".<br>The calculation will ignore sites categorized as web advertisements and content delivery networks. The browse time calculation is based on container pages logged in the URL filtering logs. Container pages are used as the basis for this calculation because many sites load content from external sites that should not be considered. For more information on the container page, refer to "Container Pages" on page 45.<br>The average browse time setting is the average time that the admin thinks it should take a user to browse a web page. Any request made after the average browse time has elapsed will be considered a new browsing activity. The calculation will ignore any new web pages that are loaded between the time of the first request (start time) and the average browse time. This behavior was designed to exclude any external sites that are loaded within the web page of interest.<br>Example: If the average browse time setting is 2 minutes and a user opens a web page and views that page for 5 minutes, the browse time for that page will still be 2 minutes. This is done because there is no way to determine how long a user views a given page.<br>(Range 0-300 seconds, default 60 seconds) |

**Table 1. Management Settings (Continued)**

| Item | Description |
|---|---|
| Page Load Threshold (sec) | Configure this variable to adjust how browse time is calculated in the "User Activity Report".<br>This option allows you to adjust the assumed time it takes for page elements to load on the page. Any request that occurs between the first page load and the page load threshold is assumed to be elements of the page. Any requests that occur outside of the page load threshold is assumed to be the user clicking a link within the page.<br>(Range 0-60 seconds, default 20 seconds) |
| Send Hostname In Syslog | Select the check box to send the device hostname field in syslog messages.<br>When this option is set, syslog messages will contain the hostname of the firewall device in their header. |
| Stop Traffic when LogDb full | Select the check box if you want traffic through the firewall to stop when the log database is full (default off). |
| Enable Log on High DP Load | Select this check box if you would like a system log entry generated if the device is under severe load (default off). |
| Buffered log forwarding from device (option in the **Panorama > Set up** tab) | Allows the firewall to buffer log entries on the device's hard disk (local storage) when it loses connectivity to Panorama. When the connection to Panorama is restored, the log entries are forwarded to Panorama; the disk space available for buffering depends on the log storage quota for the platform and the volume of logs that are pending roll over. If the available space is consumed, the oldest entries are deleted to allow logging of new events.<br>Enabled by default. |
| Get Only New Logs on Convert to Primary (option in the **Panorama > Set up** tab) | This option is only applicable when Panorama writes logs to a Network File Share (NFS). With NFS logging, only the *primary* Panorama is mounted to the NFS. Therefore, the devices send logs to the *active primary* Panorama only.<br>This option allows an administrator to configure the managed devices to only send newly generated logs to Panorama when an HA failover occurs and the secondary Panorama resumes logging to the NFS (after it is promoted as primary).<br>This behavior is typically enabled to prevent the devices from sending a large volume of buffered logs when connectivity to Panorama is restored after a significant period of time. |
| Only Active Primary Logs to Local Disk (in Panorama > Set up tab only) | Allows you to configure only the active primary Panorama to save logs to the local disk.<br>This option is valid for a Panorama virtual machine with a virtual disk and to the M-100 appliance in Panorama mode. |
| **Minimum Password Complexity** | |
| Enabled | Enable minimum password requirements for local accounts. With this feature, you can ensure that local administrator accounts on the firewall will adhere to a defined set of password requirements.<br>You can also create a password profile with a subset of these options that will override these settings and can be applied to specific accounts. For more information, refer to "Defining Password Profiles" on page 59 and refer to "Username and Password Requirements" on page 57 for information on valid characters that can be used for accounts. |

**Table 1. Management Settings (Continued)**

| Item | Description |
|---|---|
| | ***Note:*** *The maximum password length that can be entered is 31 characters. When setting requirements, make sure you do not create a combination that will not be accepted. Example, you would not be able to set a requirement of 10 uppercase, 10 lower case, 10 numbers, and 10 special characters since that would exceed the maximum length of 31.*<br><br>**Note:** *If you have High Availability (HA) configured, always use the primary device when configuring password complexity options and commit soon after making changes.* |
| Minimum Length | Require minimum length from 1-15 characters. |
| Minimum Uppercase Letters | Require a minimum number of uppercase letters from 0-15 characters. |
| Minimum Lowercase Letters | Require a minimum number of lowercase letters from 0-15 characters. |
| Minimum Numeric Letters | Require a minimum number of numeric letters from 0-15 numbers. |
| Minimum Special Characters | Require a minimum number of special characters (non-alphanumeric) from 0-15 characters. |
| Block Repeated Characters | Do no allow repeated characters based on the specified value. Example, if the value is set to 4, the password test2222 would not be accepted, but test222 would be accepted (range 2-15). |
| Block Username Inclusion (including reversed) | Select this check box to prevent the account username (or reversed version of the name) from being used in the password. |
| New Password Differs By Characters | When administrators change their passwords, the characters must differ by the specified value. |
| Require Password Change on First Login | Select this check box to prompt the administrators to change their passwords the first time they log in to the device. |
| Prevent Password Reuse Limit | Require that a previous password is not reused based on the specified count. Example, if the value is set to 4, you could not reuse the any of your last 4 passwords (range 0-50). |
| Block Password Change Period (days) | User cannot change their passwords until the specified number of days has been reached (range 0-365 days). |
| Required Password Change Period (days) | Require that administrators change their password on a regular basis specified a by the number of days set, ranging from 0-365 days. Example, if the value is set to 90, administrators will be prompted to change their password every 90 days.<br>You can also set an expiration warning from 0-30 days and specify a grace period. |
| Expiration Warning Period (days) | If a required password change period is set, this setting can be used to prompt the user to change their password at each log in as the forced password change date approaches (range 0-30 days). |

**Table 1. Management Settings (Continued)**

| Item | Description |
|---|---|
| Allowed expired admin login (count) | Allow the administrator to log in the specified number of times after the account has expired. Example, if the value is set to 3 and their account has expired, they can log in 3 more times before their account is locked out (range 0-3 logins). |
| Post Expiration Grace Period (days) | Allow the administrator to log in the specified number of days after the account has expired (range 0-30 days). |

# Defining Operations Settings

▶ *Device > Setup > Operations*

When you change a configuration setting and click **OK**, the current "candidate" configuration is updated, not the active configuration. Clicking **Commit** at the top of the page applies the candidate configuration to the active configuration, which activates all configuration changes since the last commit.

This method allows you to review the configuration before activating it. Activating multiple changes simultaneously helps avoid invalid configuration states that can occur when changes are applied in real-time.

You can save and roll back (restore) the candidate configuration as often as needed and also load, validate, import, and export configurations. Pressing **Save** creates a copy of the current candidate configuration, whereas choosing **Commit** updates the active configuration with the contents of the candidate configuration.

***Note:*** *It is a good idea to periodically save the configuration settings you have entered by clicking the* ***Save*** *link in the upper-right corner of the screen.*

To manage configurations, select the appropriate configuration management functions, as described in the following table.

**Table 2. Configuration Management Functions**

| Function | Description |
|---|---|
| **Configuration Management** | |
| Validate candidate config | Checks the candidate configuration for errors. |
| Revert to last saved config | Restores the last saved candidate configuration from the local drive. The current candidate configuration is overwritten. An error occurs if the candidate configuration has not been saved. |
| Revert to running config | Restores the last running configuration. The current running configuration is overridden. |
| Save named configuration snapshot | Saves the candidate configuration to a file. Enter a file name or select an existing file to be overwritten. Note that the current active configuration file (*running-config.xml*) cannot be overwritten. |

**Table 2. Configuration Management Functions (Continued)**

| Function | Description |
|---|---|
| Save candidate config | Saves the candidate configuration in flash memory (same as clicking **Save** at the top of the page). |
| Load named configuration snapshot | Loads a candidate configuration from the active configuration (*running-config.xml*) or from a previously imported or saved configuration. Select the configuration file to be loaded. The current candidate configuration is overwritten. |
| Load configuration version | Loads a specified version of the configuration. |
| Export named configuration snapshot | Exports the active configuration (*running-config.xml*) or a previously saved or imported configuration. Select the configuration file to be exported. You can open the file and/or save it in any network location. |
| Export configuration version | Exports a specified version of the configuration. |
| Export device state | This feature is used to export the configuration and dynamic information from a firewall that is configured as a GlobalProtect Portal with the large scale VPN feature enabled. If the Portal experiences a failure, the export file can be imported to restore the Portal's configuration and dynamic information.<br>The export contains a list of all satellite devices managed by the Portal, the running configuration at the time of the export, and all certificate information (Root CA, Server, and Satellite certificates).<br>**Important:** You must manually run the device state export or create a scheduled XML API script to export the file to a remote server. This should be done on a regular basis since satellite certificates may change often.<br>To create the device state file from the CLI, from configuration mode run `save device state`.<br>The file will be named device_state_cfg.tgz and is stored in /opt/pancfg/mgmt/device-state. The operational command to export the device state file is `scp export device-state` (you can also use `tftp export device-state`).<br>For information on using the XML API, refer to the document "PAN-OS XML-Based Rest API Usage Guide" at *https://live.paloaltonetworks.com/community/documentation*.<br>Refer to "GlobalProtect Large Scale VPN Deployment" on page 323. |
| Import named config snapshot | Imports a configuration file from any network location. Click **Browse** and select the configuration file to be imported. |
| Import device state | Import the device state information that was exported using the Export device state option. This includes the current running config, Panorama templates, and shared policies. If the device is a Global Protect Portal, the export includes the Certificate Authority (CA) information and the list of satellite devices and their authentication information. |
| **Device Operations** | |

**Table 2. Configuration Management Functions (Continued)**

| Function | Description |
|---|---|
| Reboot Device | To restart the firewall, click **Reboot Device**. You are logged out and the PAN-OS software and active configuration are reloaded. Existing sessions will also be closed and logged and a system log entry will be created that will show the administrator name that initiated the shutdown. Any configuration changes that have not been saved or committed are lost (refer to "Defining Operations Settings" on page 37).<br>***Note:*** *If the web interface is not available, use the CLI command* `request restart system`*. Refer to the PAN-OS Command Line Interface Reference Guide for details.* |
| Shutdown Device | To perform a graceful shutdown of the firewall, click **Shutdown Device** and then click **Yes** on the confirmation prompt. Any configuration changes that have not been saved or committed are lost. All administrators will be logged off and the following processes will occur:<br>• All login sessions will be logged off.<br>• Interfaces will be disabled.<br>• All system processes will be stopped.<br>• Existing sessions will be closed and logged.<br>• System Logs will be created that will show the administrator name who initiated the shutdown. If this log entry cannot be written, a warning will appear and the system will not shutdown.<br>• Disk drives will be cleanly unmounted and the device will powered off.<br>You need to unplug the power source and plug it back in before you can power on the device.<br>***Note:*** *If the web interface is not available, use the CLI command* `request shutdown system`*. Refer to the PAN-OS Command Line Interface Reference Guide for details.* |
| Restart Data Plane | To restart the data functions of the firewall without rebooting, click **Restart Dataplane**. *This option is not available on the PA-200.*<br>***Note:*** *If the web interface is not available, use the CLI command* `request restart dataplane`*. Refer to the PAN-OS Command Line Interface Reference Guide for details.* |

**Table 2. Configuration Management Functions (Continued)**

| Function | Description |
|---|---|
| **Miscellaneous** | |
| Custom Logos | Use this option to customize any of the following:<br>• Login screen background image<br>• Main UI (User Interface) header image<br>• PDF report title page image. Refer to "Managing PDF Summary Reports" on page 270.<br>• PDF report footer image<br>Click to upload an image file, to preview, or to remove a previously-uploaded image.<br>Note the following:<br>• Supported file types are png, gif, and jpg.<br>***Note:*** *Image files that contain an alpha channel are not supported and when used in PDF reports, the reports will not be generated properly. You may need to contact the illustrator who created the image to remove alpha channels in the image or make sure the graphics software you are using does not save files with the alpha channel feature.*<br>• To return to the default logo, remove your entry and commit.<br>• The maximum image size for any logo image is 128 KB.<br>• For the login screen and main user interface options, when you click , the image is shown as it will be displayed. If necessary, the image is cropped to fit. For the PDF reports, the images are auto-resized to fit without cropping. In all cases, the preview shows the recommended image dimensions.<br>For information on generating PDF reports, refer to "Managing PDF Summary Reports" on page 270. |
| SNMP Setup | Specify SNMP parameters. Refer to "SNMP" on page 47. |
| Statistics Service Setup | Specify settings for the statistics service. Refer to "Statistics Service" on page 48. |

***Note:*** *When you click* ***Commit*** *or enter a* ***commit*** *CLI command, all changes made through the web interface and the CLI since the last commit are activated. To avoid possible conflicts, use the transaction locking functions as described in "Locking Transactions" on page 27.*

# Defining Services Settings

▶ *Device > Setup > Services*

Use the **Services** tab to define settings for Domain Name System (DNS), Network Time Protocol (NTP), update servers, proxy servers, and service route configuration.

**Table 3. Services Settings**

| Function | Description |
|---|---|
| DNS | Select the type of DNS service. This setting is used for all DNS queries initiated by the firewall in support of FQDN address objects, logging, and device management. Options include:<br>• Primary and secondary DNS servers for domain name resolution<br>• DNS proxy that has been configured on the firewall |
| Primary DNS Server | Enter the IP address of the primary DNS server. The server is used for DNS queries from the firewall, for example, to find the update server, to resolve DNS entries in logs, or for FDQN-based address objects. |
| Secondary DNS Server | Enter the IP address of a secondary DNS server to use if the primary server is unavailable (optional). |
| Primary NTP Server | Enter the IP address or host name of the primary NTP server, if any. If you do not use NTP servers, you can set the device time manually. |
| Secondary NTP Server | Enter the IP address or host name of secondary NTP servers to use if the primary server is unavailable (optional). |
| Update Server | This setting represents the IP address or host name of the server used to download updates from Palo Alto Networks. The current value is **updates.paloaltonetworks.com**. Do not change the server name unless instructed by technical support. |
| **Proxy Server** | |
| Server | If the device needs to use a proxy server to reach Palo Alto Networks update services, enter the IP address or host name of the server. |
| Port | Enter the port for the proxy server. |
| User | Enter the user name to access the server. |
| Password/Confirm Password | Enter and confirm the password for the user to access the proxy server. |

**Table 3. Services Settings (Continued)**

| Function | Description |
|---|---|
| Service Route Configuration | Specify how the firewall will communicate with other servers/devices for services communication, such as DNS, Email, Palo Alto Updates, and NTP.<br>You can select "Use Management Interface for all" to use the built-in management port (MGT) for all communications, or you can define a specific source IP address to use different interfaces for granular control of each service. For example, you could configure a specific source IP of an interface for all email communication between the firewall and an email server and use a different source IP/interface for Palo Alto Updates.<br>Click **Service Route Configuration** and configure the following:<br>• **Use Management Interface for all**—This option will force all firewall service communications with external servers through the management interface (MGT). If this option is selected, you will need to configure the MGT interface to allow communications between the firewall and the servers/devices that provide services. To configure the MGT interface, navigate to Device > Setup > Management tab and edit "Management Interface Settings".<br>• **Select**—Choose this option to configure granular control for service communication. A table will appear that shows a list of available services and a drop-down to select a source address. Select the desired service and then select the source from the **Source Address** drop-down list. The "Source Address" is the IP address assigned to an interface that will be the source for the service traffic. You do not have to define the destination address since the destination is configured when configuring the given service. For example, when you define your DNS servers from the Device > Setup > Services tab and **Services**, that will set the destination for DNS queries. |

**Table 3. Services Settings (Continued)**

| Function | Description |
|---|---|
| Service Route Configuration *(Continued)* | • **Destination** and **Source Address** fields—If a service that you want to route is not listed in the **Service** column, you can use the Destination and Source Address fields to define routes that will be used by other services. Services not listed include items such as Kerberos, LDAP, and Panorama log collector communications. You do not need to enter the subnet for the destination address.<br>In multi-tenant environments, destination IP-based service routes will be required where common services require different source address. For example, if two tenants need to use RADIUS.<br>It is important that routing and policies are setup properly for the interface that will be used to route the service. For example, if you want to route Kerberos authentication requests on an interface other than the MGT port, you need to configure the Destination and Source Address in the right section of the Service Route Configuration window since Kerberos is not listed in the default Service column. As an example, you could have a source IP address 192.168.2.1 on Ethernet1/3 and then a destination for a Kerberos server of 10.0.0.240. You will need to add Ethernet1/3 to an existing virtual router with a default route, or you can create a new virtual router from Network > Virtual Routers and add static routes as needed. This will ensure that all traffic on the interface will be routed through the virtual router to reach the appropriate destinations. In this case, the destination address is 10.0.0.240 and Ethernet1/3 interface has the source IP 192.168.2.1/24.<br>The CLI output for the Destination and Source Address would look like the following:<br>`PA-200-Test# show route`<br>`destination {`<br>`10.0.0.240 {`<br>`source address 192.168.2.1/24`<br>`}`<br>With this configuration, all traffic on interface Ethernet1/3 will use the default route defined in the virtual router and will be sent to 10.0.0.240. |

# Defining Content ID Settings

▶ ***Device > Setup > Content-ID***

Use the **Content-ID** tab to define settings for URL filtering, data protection, and container pages.

**Table 4. Content ID Settings**

| Function | Description |
|---|---|
| **URL Filtering** | |

**Table 4. Content ID Settings (Continued)**

| Function | Description |
|---|---|
| Dynamic URL Cache Timeout | Click **Edit** and enter the timeout (in hours). This value is used in dynamic URL filtering to determine the length of time an entry remains in the cache after it is returned from the URL filtering service. This option is applicable to URL filtering using the BrightCloud database only. For information on URL filtering, refer to "URL Filtering Profiles" on page 217. |
| URL Continue Timeout | Specify the interval following a user's "continue" action before the user must press continue again for URLs in the same category (range 1 - 86400 minutes, default 15 minutes). |
| URL Admin Override Timeout | Specify the interval after the user enters the admin override password before the user must re-enter the admin override password for URLs in the same category (range 1 - 86400 minutes, default 900 minutes). |
| URL Admin Lockout Timeout | Specify the period of time that a user is locked out from attempting to use the URL Admin Override password following three unsuccessful attempts (1 - 86400 minutes, default 1800 minutes). |
| x-forwarded-for | Include the X-Forwarded-For header that includes the source IP address. When this option is selected, the firewall examines the HTTP headers for the X-Forwarded-For header, which a proxy can use to store the original user's source IP address.<br>The system takes the value and places **Src: x.x.x.x** into the Source User field of the URL logs (where *x.x.x.x* is the IP address that is read from the header). |
| Strip-x-forwarded-for | Remove the X-Forwarded-For header that includes the source IP address. When this option is selected, the firewall zeros out the header value before forwarding the request, and the forwarded packets do not contain internal source IP information. |
| Allow Forwarding of Decrypted Content | Select the check box to allows the firewall to forward decrypted content to an outside service. For example, when this option is set the firewall can send decrypted content to WildFire for analysis. For multi-VSYS configurations, this option is per VSYS. |
| **URL Admin Override** | |
| Settings for URL Admin Override | Specify the settings that are used when a page is blocked by the URL filtering profile and the **Override** action is specified. Refer to "URL Filtering Profiles" on page 217.<br>Click **Add** and configure the following settings for each virtual system that you want to configure for URL admin override.<br>• **Location**—Select the virtual system from the drop-down list (multi-VSYS devices only).<br>• **Password/Confirm Password**—Enter the password that the user must enter to override the block page.<br>• **Server Certificate**—Select the server certificate to be used with SSL communications when redirecting through the specified server.<br>• **Mode**—Determines whether the block page is delivered transparently (it appears to originate at the blocked website) or redirected to the user to the specified server. If you choose **Redirect**, enter the IP address for redirection.<br>Click ☒ to delete an entry. |

**Table 4. Content ID Settings (Continued)**

| Function | Description |
|---|---|
| Manage Data Protection | Add additional protection for access to logs that may contain sensitive information, such as credit card numbers or social security numbers.<br>Click **Manage Data Protection** and configure the following:<br>• To set a new password if one has not already been set, click **Set Password**. Enter and confirm the password.<br>• To change the password, click **Change Password**. Enter the old password, and enter and confirm the new password.<br>• To delete the password and the data that has been protected, click **Delete Password**. |
| Container Pages | Use these settings to specify the types of URLs that the firewall will track or log based on content type, such as application/pdf, application/soap+xml, application/xhtml+, text/html, text/plain, and text/xml. Container pages are set per virtual system, which you select from the **Location** drop-down list. If a virtual system does not have an explicit container page defined, the default content types are used.<br>Click **Add** and enter or select a content type.<br>Adding new content types for a virtual system overrides the default list of content types. If there are no content types associated with a virtual system, the default list of content types is used. |

# Defining Session Settings

▶ *Device > Setup > Session*

The **Session** tab allows you to configure session age-out times and global session-related settings such as firewalling IPv6 traffic and rematching security policy to existing sessions when the policy changes.

**Table 5. Session Settings**

| Field | Description |
|---|---|
| **Session Settings** | |
| Rematch Sessions | Click **Edit** and select the check box **Rematch all sessions on config policy change**.<br>For example, assume that Telnet was previously allowed and then changed to **Deny** in the last commit. The default behavior is for any Telnet sessions that were started before the commit to be rematched and blocked. |
| ICMPv6 Token Bucket Size | Enter the bucket size for rate limiting of ICMPv6 error messages. The token bucket size is a parameter of the token bucket algorithm that controls how bursty the ICMPv6 error packets can be (range 10-65535 packets, default 100). |
| ICMPv6 Error Packet Rate | Enter the average number of ICMPv6 error packets per second allowed globally (range 10-65535 packets/sec, default 100). This value applies to all interfaces. |
| Jumbo Frame<br>Jumbo Frame MTU | Select to enable jumbo frame support. The Default value for Jumbo frames is 9192. (range 512 - 9216 bytes). Refer to the spec sheet for your firewall model, available at *http://www.paloaltonetworks.com*. |

**Table 5. Session Settings (Continued)**

| Field | Description |
|---|---|
| Enable IPv6 Firewalling | To enable firewall capabilities for IPv6, click **Edit** and select the **IPv6 Firewalling** check box.<br>All IPv6-based configurations are ignored if IPv6 is not enabled. |
| NAT64 IPv6 Minimum Network MTU | Allows you to change the global MTU setting for IPv6 translated traffic (default 1280). The default is based on the standard minimum MTU for IPv6 traffic. |
| Accelerated Aging | Allows for the accelerated aging-out of idle sessions.<br>Select the check box to enable accelerated aging and specify the threshold (%) and scaling factor.<br>When the session table reaches the **Accelerated Aging Threshold** (% full), the **Accelerated Aging Scaling Factor** is applied to the aging calculations for all sessions. The session's idle time is calculated as the actual idle time multiplied by the scaling factor. For example, if a scaling factor of 10 is used, a session that would normally time out after 3600 seconds would instead time out 10 times faster, so it would time out in 360 seconds. |
| **Session Timeouts** | |
| Timeouts | Specify timeouts in seconds for each of the categories. Ranges and defaults are listed. |
| **Session Features** | |
| Decryption Certificate Revocation Settings | Select to configure the certificate management options for the firewall. |
| Enable | Select the check box to use Certificate Revocation Lists (CRL) to check the validity of SSL certificates.<br>Each trusted certificate authority (CA) maintains CRL to determine if an SSL certificate is valid (not revoked) for SSL decryption. The Online Certificate Status Protocol (OCSP) can also be used to dynamically check the revocation status of a certificate. For more information on SSL decryption, refer to "Decryption Policies" on page 202. |
| Receive Timeout | Specify the interval after which the CRL request times out and the status is determined to be unknown (1-60 seconds). |
| Enable OCSP | Select the check box to use OCSP to check the validity of SSL certificates. |
| Receive Timeout | Specify the interval after which the OCSP requests times out and the status is determined to be unknown (1-60 seconds). |
| Block Session With Unknown Certificate Status | Select the check box if you want to block certificates that cannot be validated. |
| Block Session On Certificate Status Check Timeout | Select the check box if you want to block certificates when the request for certificate information times out. |
| Certificate Status Timeout | Specify the interval after which certificate status requests time out (1-60 seconds). |

# SNMP

▶ *Device > Setup > Operations*

Use this page to define access to SNMP Management Information Bases (MIBs) for SNMPv2c and SNMPv3. Click **SNMP Setup** on the **Setup** page, and specify the following settings.

A MIB module defines all SNMP traps generated by the system. Each event log in the system is defined as an independent SNMP trap with an Object ID (OID) of its own, and individual fields in an event log are defined as a variable binding (varbind) list.

**Table 6. SNMP Setup**

| Field | Description |
|---|---|
| Physical Location | Specify the physical location of the firewall. |
| Contact | Enter the name or email address of the person responsible for maintaining the firewall. This setting is reported in the standard system information MIB. |
| Use Specific Trap Definitions | Select the check box to use a unique OID for each SNMP trap based on the event type (default is selected). |
| Version | Select the SNMP version (V2c or V3). This setting controls access to the MIB information. By default, V2c is selected with the "public" community string.<br>For V2c, configure the following setting:<br>• **SNMP Community String**—Enter the SNMP community string for firewall access (default **public**).<br><br>For V3, configure the following settings:<br>• **Views**—Click **Add** and configure the following settings:<br>– **Name**—Specify a name for a group of views.<br>– **View**—Specify a name for a view.<br>– **OID**—Specify the object identifier (OID) (for example, 1.2.3.4).<br>– **Option**—Choose whether the OID is to be included or excluded from the view.<br>– **Mask**—Specify a mask value for a filter on the OID in hexadecimal format (for example, 0xf0).<br>• **Users**—Click **Add** and configure the following settings:<br>– **Users**—Specify a user name.<br>– **View**—Specify the group of views for the user.<br>– **Auth Password**—Specify the user's authentication password (minimum 8 characters, maximum of 256 characters, and no character restrictions). All characters allowed). Only Secure Hash Algorithm (SHA) is supported.<br>– **Priv Password**—Specify the user's encryption password (minimum 8 characters, maximum of 256 characters, and no character restrictions). Only Advanced Encryption Standard (AES) is supported. |

# Statistics Service

▶ *Device > Setup > Operations*

The **Statistics Service** feature allows the firewall to send anonymous application, threat, and crash information to the Palo Alto Networks research team. The information collected enables the research team to continually improve the effectiveness of Palo Alto Networks products based on real-world information. This service is disabled by default and once enabled, information will be uploaded every 4 hours.

You can allow the firewall to send any of the following types of information:

- Application and Threat Reports
- Unknown Application Reports
- URL Reports
- Device traces for crashes

To view a sample of the content for a statistical report to be sent, click the report icon . The **Report Sample** tab opens to display the report code. To view a report, click the check box next to the desired report, then click the **Report Sample** tab.

# Comparing Configuration Files

▶*Device > Config Audit*

You can view and compare configuration files by using the **Config Audit** page. From the drop-down lists, select the configurations to compare. Select the number of lines that you want to include for context, and click **Go**.

The system presents the configurations and highlights the differences, as in the following figure.

The page also includes << and >> buttons adjacent to the drop-down lists, which are enabled when comparing two consecutive configuration versions. Click << to change the configurations being compared to the previous set of stored configurations, and click to >> to change the configurations being compared to the next set of stored configurations.

**Figure 1. Configuration Comparison**

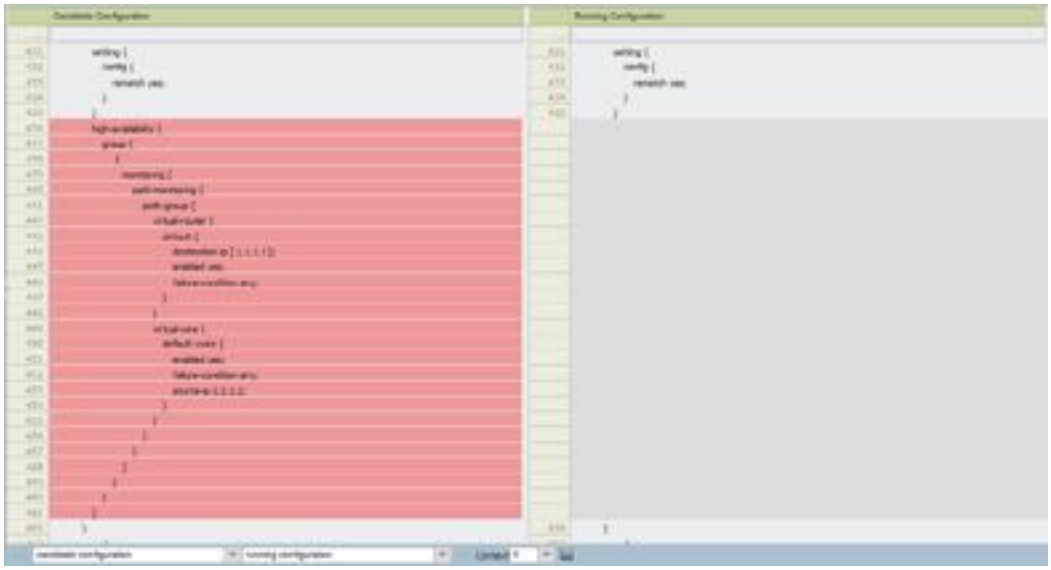

Panorama automatically saves all of the configuration files that are committed on each managed firewall, whether the changes are made through the Panorama interface or locally on the firewall.

# Installing a License

▶ *Device > Licenses*

When you purchase a subscription from Palo Alto Networks, you receive an authorization code to activate one or more license keys.

To activate a URL filtering license, you must install the license, download the database, and click **Activate**.

The following functions are available on the **Licenses** page:

- To enable licenses for URL filtering, click **Activate**. If you are using PAN-DB for URL Filtering, you will need to click **Download** to retrieve the initial seed database first and then click Activate. You can also run the CLI request url-filtering download paloaltonetworks region <region name>.
- To enable purchased subscriptions that require an authorization code and have been activated on the support portal, click **Retrieve license keys from license server**.
- To enable purchased subscriptions that require an authorization code and have not been previously activated on the support portal, click **Activate feature using authorization code**. Enter your authorization code, and click **OK**.
- If the firewall does not have connectivity to the license server and you want to upload license keys manually, follow these steps:
  a. Obtain a file of license keys from *http://support.paloaltonetworks.com*.
  b. Save the license key file locally.
  c. Click **Manually upload license key**, click **Browse** and select the file, and click **OK**.

### Important items to consider when installing a license

If you are unable to activate the URL filter using the web interface, CLI commands are available. Refer to the *PAN-OS Command Line Interface Reference Guide* for more information.

# Upgrading/Downgrading the PAN-OS Software

▶ *Device > Software*

To upgrade to a new release of the PAN-OS software, you can view the latest versions of the PAN-OS software available from Palo Alto Networks, read the release notes for each version, and then select the release you want to download and install (a support license is required).

Perform any of the following functions on the **Software** page:

- Click **Check Now** to retrieve the latest software releases available from Palo Alto Networks.
- Click **Release Notes** to view a description of the changes in a release and to view the migration path to install the software.
  You cannot skip a feature release and must have a base image downloaded before upgrading from one feature release to a maintenance release in a higher feature release. This is due to the fact that the maintenance release does not contain the full software

build, only changes made since the base image release, so the base image is needed for the upgrade. For example, if you upgrade from 4.0.12 to 4.1.7, you need to download the 4.1.0 base image (not install), so the 4.1.7 maintenance release upgrade can access the 4.1.0 base image files. If you want to upgrade to from one release to two higher feature releases, you need to upgrade to each feature release. For example, if you want to upgrade from 4.0 to 5.0, you have to upgrade from 4.0 to 4.1, then 4.1 to 5.0.
You can delete the base image files when you complete the upgrade, but that is not recommended since the base image would be needed when you upgrade to the next maintenance release. You would only want to delete the base image for older releases that will not need upgrading. For example, if you are running 4.1, you probably do not need the base images for 3.1 and 4.0 unless you plan on downgrading to those versions.

- Click **Download** to install a new release from the download site. When the download is complete, a checkmark is displayed in the **Downloaded** column. To install a downloaded release, click **Install** next to the release.

  During installation, you are asked whether to reboot when installation is complete. When the installation is complete, you will be logged out while the firewall is restarted. The firewall will be rebooted, if that option was selected.

  When upgrading to a feature release (where the first or second digit in the PAN-OS version changes, e.g. 4.0 to 4.1. or 4.1 to 5.0), you will see a message when you click install that indicates that you are about to perform a feature release upgrade. You should make sure you back up your current configuration since a feature release may migrate certain configurations to accommodate new features. Refer to "Downgrading PAN-OS Software" on page 53.

- Click **Upload** to install a release that you previously stored on your PC. Browse to select the software package, and click **Install from File**. Choose the file that you just selected from the drop-down list, and click **OK** to install the image.

- Click the Delete icon to delete an outdated release.

### Items to note when upgrading the PAN-OS software

- When upgrading from an earlier PAN-OS version, follow the recommended path to reach the latest release, as described in the release notes.

- When upgrading a High Availability (HA) pair to a new feature release (where the first or second digit in the PAN-OS version changes, e.g. 4.0 to 4.1. or 4.1 to 5.0), the configuration may be migrated to accommodate new features. If session synchronization is enabled, sessions will not be synchronized if one device in the cluster is at a different PAN-OS feature release.

- The date and time settings on the firewall must be current. PAN-OS software is digitally signed and the signature is checked by the device prior to installing a new version. If the date setting on the firewall is not current, the device may perceive the software signature to be erroneously in the future and will display the message
  `Decrypt failed: GnuPG edit non-zero, with code 171072 Failed to load into PAN software manager.`

# Upgrading PAN-OS in a High Availability Configuration

This section provides information on upgrading the PAN-OS software on a high availability (HA) configuration in active/passive and active/active modes.

If session synchronization is enabled, sessions will not be synchronized if one device in the cluster is at a different PAN-OS feature release. If session continuity is required, you must temporarily permit non-syn-tcp while the session table is rebuilt.

For information on High Availability (HA) configurations, refer to "High Availability" on page 93. To view the general status of your HA configuration, go to the **Dashboard** tab and view the **High Availability** widget. If the widget is not displayed, click the Widgets link drop-down and select **System > High Availability**.

The following steps assume that you are performing a new feature release upgrade. The same procedures can be used for a maintenance release, but in that case if session synchronization is enabled, there is no interruption to that process while both devices are functional. Also, the higher revision maintenance release in the pair will not go non-functional as it does with a feature release.

To upgrade an HA pair to a new feature release:

***Note:*** *In the following steps, active/passive and active/active applies, the terminology is just different. In active/passive you have an active and a passive device. In active/active you have an active-primary and active-secondary and both devices are passing traffic when the cluster is functional.*

1. Read "Upgrading/Downgrading the PAN-OS Software" on page 50 for general information and important notes on the upgrade process.
2. Save a backup of the current configuration file on both devices by navigating to **Device > Setup > Operations** tab and select **Export named configuration snapshot**, select **running-config.xml** and then click **OK** to save the configuration file to your computer.
3. Upgrade the passive/active-secondary firewall (device B) to the new PAN-OS feature release and then reboot to complete the installation. After the reboot, the device will not be functional until device A (active/active-primary) is suspended.
   If session synchronization is enabled, you can run the operational command `set session tcp-reject-non-syn no`. This will rebuild the session table and sessions prior to the upgrade will continue. Once both devices are at the same feature release, run `set session tcp-reject-non-syn yes` to enable this option.
   In active/passive mode, if the preemptive option is configured, the current passive device will revert to active when state synchronization is complete.
4. Suspend the active/active-primary firewall (device A), which will force the passive/active-secondary (device B) firewall to become functional. At this point, session synchronization will stop due to the fact that both devices are not on the same PAN-OS feature release. Make sure device B is functional and passing traffic before continuing the upgrade on device A.
5. Upgrade "device A" to the new PAN-OS release and then reboot the device to complete the installation. When device A comes back up, all HA functions should resume, including session synchronization.
   If you are allowing non-syn-tcp as indicated in Step 3, you can revert back by running `set session tcp-reject-non-syn yes`.
6. Verify that the active device is passing traffic by viewing the **Monitor > Session Browser**, or by running `show session all` from the CLI.
   You can also check the state of HA on the device by running `show high-availability all | match reason`. If this is an active/active configuration, check that both devices are passing traffic.
   To check session synchronization run `show high-availability interface ha2`. In the

"Hardware Interface counters read from CPU" table check that counters are increasing. In an active/passive configuration, only the active device will show packets transmitted and the passive device will only show packets received. In active/active, you will see packets received and packets transmitted on both devices.

# Downgrading PAN-OS Software

If you need to downgrade the PAN-OS software, you will need to follow different steps based on the PAN-OS release type (feature or maintenance). In a feature release (where the first or second digit in the PAN-OS version changes, e.g. 4.0 to 4.1. or 4.1 to 5.0), the configuration may be migrated to accommodate new features, so you should not downgrade unless you also restore the configuration for the previous release. Maintenance releases can be downgraded without having to worry about restoring the configuration.

- "Maintenance Release Downgrade" in the next section
- "Feature release Downgrade" on page 54

**CAUTION:** It is recommended that you downgrade into a configuration that matches the software version. Unmatched software and configurations can result in failed downgrades or even force the system into maintenance mode. This only applies to a downgrade from one feature release to another, not maintenance releases.

If you have a problem with a downgrade, you may need to enter maintenance mode and reset the device to factory default and then restore the configuration from the original config file that was exported prior to the upgrade.

## Maintenance Release Downgrade

A maintenance release is where the third digit in the release numbers changes (4.1.4 to 4.1.5 or 5.0.0 to 5.0.1). When you upgrade from one maintenance release to another, there are no feature changes, so there is no need to restore the configuration during the downgrade.

To downgrade to an earlier maintenance release:

1. Save a backup of the current configuration file by navigating to the **Device > Setup > Operations** tab and select **Export named configuration snapshot**, select **running-config.xml** and then click **OK** to save the configuration file. This backup can be used to restore the configuration if you have problems with the downgrade and you need to do a factory reset.
2. Navigate to **Device > Software** and you will see the software page that lists all PAN-OS versions that can be downloaded, or that have already been downloaded.
3. To downgrade to an older maintenance release, click **Install** in the **Action** column for the desired release. If the version you want to use shows **Download**, click the **Download** link to retrieve the software package and then click **Install**.
4. After PAN-OS has been downgraded, click **OK** to reboot the device to activate the new version.

## Feature release Downgrade

A feature release is where the first or second digit in the release numbers changes (4.0 to 4.1, or 4.1 to 5.0). When you upgrade from one feature release to another, the configuration may be changed to accommodate the new features. Due to this fact, you will need to restore your configuration backup that was created before the device was upgraded to a feature release. As of PAN-OS 4.1, this configuration backup is automatically created when upgrading to a feature release.

To downgrade to an earlier feature release:

**Note:** It is important to note that this procedure will restore your device to the configuration that was in place before the upgrade to a feature release. Any changes made since that time will be lost, so it is important to back up your current configuration in case you want to restore those changes when you return to the newer release.

1. Save a backup of the current configuration file by navigating to the **Device > Setup > Operations** tab and select **Export named configuration snapshot**, select **running-config.xml** and then click **OK** to save the configuration file. This backup can be used to restore the configuration if you have problems with the downgrade and you need to do a factory reset.
2. To downgrade to an older feature release, you must locate the previous feature release by navigating to **Device > Software** and browsing to the page that contains the release. See Figure 2.
3. Click **Install** in the **Action** column for the feature release. If the version you want to use shows **Download**, click the **Download** link to retrieve the software package and then click **Install**.
4. You will see a prompt that allows you to select a configuration that will be used after the device reboots. In most cases, since this is a feature release downgrade, you want to select the auto saved configuration that was created when the device was upgraded to the next feature release. For example, if you are running PAN-OS 5.0 and want to downgrade to

PAN-OS 4.1, click **Install** in the **Action** column, then in the **Select a Config File for Downgrading** drop-down, select *autosave-4.1.0* as shown in Figure 2.

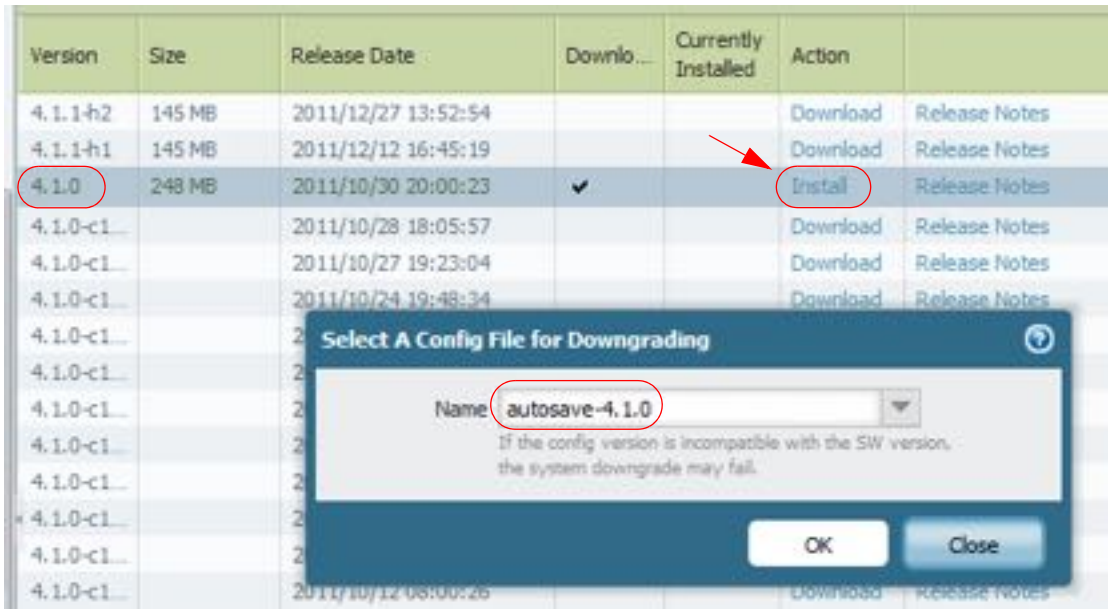


**Figure 2. Downgrade to an older feature release**

5. After PAN-OS has been installed, click **OK** to reboot the device to activate the new version. The selected feature release of PAN-OS and the configuration will then be activated.

***Note:*** *For releases prior to 4.1, you may need to do a factory reset and restore the device. For example, downgrading from 4.0 to 3.1.*

# Updating Threat and Application Definitions

▶ *Device > Dynamic Updates*

Palo Alto Networks periodically posts updates with new or revised application definitions, information on new security threats, such as antivirus signatures (threat prevention license required), URL filtering criteria, updates to GlobalProtect data, and WildFire signatures (WildFire subscription required). You can view the latest updates, read the release notes for each update, and then select the update you want to download and install. You can also revert to a previously installed version of an update.

Perform any of the following functions on this page:

- Click **Check Now** to obtain the latest information from Palo Alto Networks.
- Click **Install** in the **Action** column to update to that version.
- Click **Revert** for a version to return to that version.

- Click **Release Notes** to view a description of an update.
- Click **Upload** to install a file that you previously stored on your PC. Browse to select the file, and click **Install from File**. Choose the file that you just selected from the drop-down list, and click **OK** to install.
- Click the **Schedule** link to schedule automatic updates. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select **Download Only**, you can install the downloaded update by clicking the **Install** link in the **Action** column on the **Dynamic Updates** page. When you click **OK**, the update is scheduled. No commit is required. You can also indicate how persistent the content must be (number of hours) for the action to take place and whether the upload should be synchronized to peer firewalls.

# Administrator Roles, Profiles, and Accounts

The firewall supports the following options to authenticate administrative users who attempt to log in to the firewall:

- **Local database**—The user login and password information is entered directly into the firewall database.
- **RADIUS**—Existing Remote Authentication Dial In User Service (RADIUS) servers are used to authenticate users.
- **LDAP**—Existing Lightweight Directory Access Protocol (LDAP) servers are used to authenticate users.
- **Kerberos**—Existing Kerberos servers are used to authenticate users.
- **Client Certificate**—Existing client certificates are used to authenticate users.

When you create an administrative account, you specify local authentication or client certificate (no authentication profile), or an authentication profile (RADIUS, LDAP, Kerberos, or local DB authentication). This setting determines how the administrator's authentication is checked.

Administrator roles determine the functions that the administrator is permitted to perform after logging in. You can assign roles directly to an administrator account, or define role profiles, which specify detailed privileges, and assign those to administrator accounts.

Refer to the following sections for additional information:

- For instructions on setting up authentication profiles, refer to "Setting Up Authentication Profiles" on page 62.
- For instructions on setting up role profiles, refer to "Defining Administrator Roles" on page 58.
- For instructions on setting up administrator accounts, refer to "Creating Administrative Accounts" on page 59.
- For information on SSL virtual private networks (VPNs), refer to "Configuring GlobalProtect" on page 335.
- For instructions on defining virtual system domains for administrators, refer to "Specifying Access Domains for Administrators" on page 61.

- For instructions on defining certificate profiles for administrators, refer to "Certificate Profile" on page 89.

## Username and Password Requirements

The following table lists the valid characters that can be used in usernames and passwords for PAN-OS and Panorama accounts.

**Note:** The following restrictions apply to PAN-OS 3.1 and later.

**Table 7 Valid Characters for Usernames and Passwords**

| Account Type | Restrictions |
|---|---|
| Password Character Set | There are no restrictions on any password field character sets. |
| Remote Admin, SSL-VPN, or Captive Portal | The following characters are not allowed for the username:<br>• Backtick (`)<br>• Angular brackets (< and >)<br>• Ampersand (&)<br>• Asterisk (*)<br>• At sign (@)<br>• Question mark (?)<br>• Pipe (\|)<br>• Single-Quote (')<br>• Semicolon (;)<br>• Double-Quote (")<br>• Dollar ($)<br>• Parentheses ( '(' and ')' )<br>• Colon (':') |
| Local Administrator Accounts | The following are the allowed characters for local usernames:<br>• Lowercase (a-z)<br>• Uppercase (A-Z)<br>• Numeric (0-9)<br>• Underscore (_)<br>• Period (.)<br>• Hyphen (-)<br>***Note:*** *Login names cannot start with a hyphen (-).* |

***Note:*** *The system does not impose restrictions on the response page character set.*

## Defining Administrator Roles

▶ *Device > Admin Roles*

Use the **Admin Roles** page to define role profiles that determine the access and responsibilities available to administrative users. For instructions on adding administrator accounts, refer to "Creating Administrative Accounts" on page 59.

There are also three pre-defined Admin Roles that can be used for common criteria purposes. You first use the Superuser role for the initial configuration of the device and to create the administrator accounts for the Security Administrator, Audit Administrator, and Cryptographic Administrator. Once the accounts are created and the proper common criteria Admin Roles are applied, you then login using those accounts. The default Superuser account in FIPS or CC mode is **admin** and has a default password of **paloalto**. In standard operating mode, the default **admin** password is **admin**. The pre-defined Admin Roles were created where there is no overlap in capabilities, except that all have read-only access to the audit trail (except audit administrator with full read/delete access. These admin roles cannot be modified and are defined as follows:

- auditadmin—The Audit Administrator is responsible for the regular review of the firewall's audit data.
- cryptoadmin—The Cryptographic Administrator is responsible for the configuration and maintenance of cryptographic elements related to the establishment of secure connections to the firewall.
- securityadmin—The Security Administrator is responsible for all other administrative tasks (e.g. creating the firewall's security policy) not addressed by the other two administrative roles.

**Table 8. Administrator Role Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify this administrator role (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter an optional description of the role (up to 255 characters). |
| Role | Select the general scope of administrative responsibility from the drop-down list. |
| WebUI | Click the icons for specified areas to indicate the type of access permitted for the web interface:<br>• Read/write access to the indicated page.<br>• Read only access to the indicated page.<br>• No access to the indicated page. |
| CLI Role | Select the type of role for CLI access:<br>• **disable**—Access to the device CLI not permitted.<br>• **superuser**—Full access to the current device.<br>• **superreader**—Read-only access to the current device.<br>• **deviceadmin**—Full access to a selected device, except for defining new accounts or virtual systems.<br>• **devicereader**—Read-only access to a selected device. |

## Defining Password Profiles

▶ *Device > Password Profiles*

Password profiles allow you to set basic password requirements for an individual local account. If you enable **Minimum Password Complexity**, which provides password requirements for all local accounts, this password profile will override those settings. Refer to "Minimum Password Complexity" on page 35 for more information and refer to "Username and Password Requirements" on page 57 for information on valid characters that can be used for accounts.

To apply a password profile to an account, go to **Device > Administrators**, select an account and then choose the profile from the **Password Profile** drop-down.

**Table 9 Password Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the password profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Required Password Change Period (days) | Require that administrators change their password on a regular basis specified a by the number of days set, ranging from 0-365 days. Example, if the value is set to 90, administrators will be prompted to change their password every 90 days.<br>You can also set an expiration warning from 0-30 days and specify a grace period. |
| Expiration Warning Period (days) | If a required password change period is set, this setting can be used to prompt the user to change their password at each log in as the forced password change date approaches (range 0-30 days). |
| Post Expiration Admin Login Count | Allow the administrator to log in the specified number of times after their account has expired. Example, if the value is set to 3 and their account has expired, they can log in 3 more times before their account is locked out (range 0-3 logins). |
| Post Expiration Grace Period (days) | Allow the administrator to log in the specified number of days after their account has expired (range is 0-30 days). |

## Creating Administrative Accounts

▶ *Device > Administrators*

Administrator accounts control access to the firewall. Each administrator can have full or read-only access to a single device or to a virtual system on a single device. The predefined **admin** account has full access.

The following authentication options are supported:

- Password authentication—The user enters a user name and password to log in. No certificates are required.
- Client certificate authentication (web)—If you select this check box, a user name and password are not required; the certificate is sufficient to authenticate access to the firewall.

- Public key authentication (SSH)—The user can generate a public/private key pair on the machine that requires access to the firewall, and then upload the public key to the firewall to allow secure access without requiring the user enter a user name and password.

***Note:*** *To ensure that the device management interface remains secure, it is recommended that administrative passwords be changed periodically using a mixture of lower-case letters, upper-case letters, and numbers. You can also enforce "Minimum Password Complexity" from Setup > Management.*

**Table 10. Administrator Account Settings**

| Field | Description |
|---|---|
| Name | Enter a login name for the user (up to 15 characters). The name is case sensitive and must be unique. Use only letters, numbers, hyphens, periods, and underscores.<br>***Note:*** *Login names cannot start with a hyphen (-).* |
| Authentication Profile | Select an authentication profile for administrator authentication according to the settings in the specified authentication profile. This setting can be used for RADIUS, LDAP, Kerberos, or Local DB authentication.<br>For instructions on setting up authentication profiles, refer to "Setting Up Authentication Profiles" on page 62. |
| Use only client certificate authentication (web) | Select the check box to use client certificate authentication for web access. If you select this check box, a user name and password are not required; the certificate is sufficient to authenticate access to the firewall. |
| New Password<br>Confirm New Password | Enter and confirm a case-sensitive password for the user (up to 15 characters). *You can also enforce "Minimum Password" from Setup > Management.* |
| Use Public Key Authentication (SSH) | Select the check box to use SSH public key authentication. Click **Import Key** and browse to select the public key file. The uploaded key is displayed in the read-only text area.<br>Supported key file formats are IETF SECSH and OpenSSH. Supported key algorithms are DSA (1024 bits) and RSA (768-4096 bits).<br>***Note:*** *If the public key authentication fails, a username and password prompt is presented to the user.* |

**Table 10. Administrator Account Settings (Continued)**

| Field | Description |
|---|---|
| Role | Select an option for assigning a role to this user. The role determines what the user can view and modify.<br>If you choose **Dynamic**, you can select any of the following pre-specified roles from the drop-down list:<br>• **Superuser**—Full access to the current device.<br>• **Superuser (read-only)**—Read-only access to the current device.<br>• **Device Admin**—Full access to a selected device, except for defining new accounts or virtual systems.<br>• **Device administrator (read-only)**—Read-only access to a selected device.<br>• **Vsys Admin**—Full access to a selected virtual system on a specific device (if multiple virtual systems are enabled).<br>• **Vsys Admin (read-only)**—Read-only access to a selected virtual system on a specific device.<br>• **Role Based Admin**—Access based on assigned roles, as defined in "Defining Administrator Roles" on page 58.<br>If you choose **Role Based**, select a previously-defined role profile from the drop-down list. For instructions on defining role profiles, refer to "Defining Administrator Roles" on page 58. |
| Virtual System | Select the virtual systems that you want the administrator to have access to, and click **Add** to move them from the **Available** area to the **Selected** area. |

***Note:*** *On the Panorama* ***Administrators*** *page for "superuser," a lock icon is shown in the right column if an account is locked out. The administrator can click the icon to unlock the account.*

## Specifying Access Domains for Administrators

▶ *Device > Access Domain*

Use the **Access Domain** page to specify domains for administrator access to the firewall. The access domain is linked to RADIUS Vendor-specific Attributes (VSAs) and is supported only if a RADIUS server is used for administrator authentication. For information on configuring RADIUS, refer to "Configuring RADIUS Server Settings" on page 65.

When an administrator attempts to log in to the firewall, the firewall queries the RADIUS server for the administrator's access domain. If there is an associated domain on the RADIUS server, it is returned and the administrator is restricted to the defined virtual systems inside

the named access domain on the device. If RADIUS is not used, the access domain settings on this page are ignored. For information on Panorama Access Domains, refer to "Specifying Panorama Access Domains for Administrators" on page 392.

**Table 11. Access Domain Settings**

| Field | Description |
|---|---|
| Name | Enter a name for the access domain (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, hyphens, and underscores. |
| Virtual Systems | Select virtual systems in the Available column and click **Add** to select them.<br>***Note:*** *Access Domains are only supported on devices that support virtual systems.* |

# Authentication Profiles

Authentication profiles specify local database, RADIUS, LDAP, or Kerberos settings and can be assigned to administrator accounts, SSL-VPN access, and captive portal. When an administrator attempts to log in to the firewall directly or through an SSL-VPN or captive portal, the firewall checks the authentication profile that is assigned to the account and authenticates the user based on the authentication settings.

If the user does not have a local administrator account, the authentication profile that is specified on the device **Setup** page determines how the user is authenticated (refer to "Defining Management Settings" on page 30):

- If you specify **RADIUS** authentication settings on the **Setup** page and the user does not have a local account on the firewall, then the firewall requests authentication information for the user (including role) from the RADIUS server. The Palo Alto Networks RADIUS dictionary file containing the attributes for the various roles is available on the support site at *https://live.paloaltonetworks.com/docs/DOC-3189.*
- If **None** is specified as the authentication profile on the **Settings** page, then the user must be authenticated locally by the firewall according to the authentication profile that is specified for the user.

## Setting Up Authentication Profiles

▶ ***Device > Authentication Profile***

Use the **Authentication Profile** page to configure authentication settings that can be applied to accounts to manage access to the firewall.

**Table 12. Authentication Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |

**Table 12. Authentication Profile Settings (Continued)**

| Field | Description |
|---|---|
| Lockout Time | Enter the number of minutes that a user is locked out if the number of failed attempts is reached (0-60 minutes, default 0). 0 means that the lockout is in effect until it is manually unlocked. |
| Failed Attempts | Enter the number of failed login attempts that are allowed before the account is locked out (1-10, default 0). 0 means that there is no limit. |
| Allow List | Specify the users and groups that are explicitly allowed to authenticate. Click **Edit Allow List** and do any of the following:<br>• Select the check box next to the appropriate user or user group in the **Available** column, and click **Add** to add your selections to the **Selected** column.<br>• Use the **All** check box to apply to all users.<br>• Enter the first few characters of a name in the **Search** field to list all the users and user groups that start with those characters. Selecting an item in the list sets the check box in the **Available** column. Repeat this process as often as needed, and then click **Add**.<br>• To remove users or user groups, select the appropriate check boxes in the **Selected** column and click **Remove**, or select **any** to clear all users. |
| Authentication | Choose the type of authentication:<br>• **None**—Do not use any authentication on the firewall.<br>• **Local Database**—Use the authentication database on the firewall.<br>• **RADIUS**—Use a RADIUS server for authentication.<br>• **LDAP**—Use LDAP as the authentication method.<br>• **Kerberos**—Use Kerberos as the authentication method. |
| Server Profile | If you select RADIUS, LDAP, or Kerberos as the authentication method, choose the authentication server from the drop-down list. Servers are configured on the **Server** pages. Refer to "Configuring RADIUS Server Settings" on page 65, "Configuring LDAP Server Settings" on page 66, and "Configuring Kerberos Settings (Native Active Directory Authentication)" on page 67. |
| Login Attribute | If you selected LDAP as the authentication method, enter the LDAP directory attribute that uniquely identifies the user. |

**Table 12. Authentication Profile Settings (Continued)**

| Field | Description |
|---|---|
| Password Expiry Warning | If you selected LDAP as the authentication method, enter the number of days prior to password expiration to send an automated message to the user. If the field is left blank, no warning is provided. This is supported for the following databases: Active Directory, eDirectory, and Sun ONE Directory.<br>This setting is used for SSL-VPN. For more information, refer to "Configuring GlobalProtect" on page 335.<br>You can customize the expiration warning message as part of the SSL-VPN login page by editing the script<br>`<SCRIPT>`<br>`function getPassWarnHTML(expdays)`<br>`{`<br>`var str = "Your password will expire in " + expdays + " days";`<br>`return str;`<br>`}`<br>`</SCRIPT>`<br>Changing the value of the str variable changes the displayed message. |

## Creating a Local User Database

You can set up a local database on the firewall to store authentication information for remote access users, device administrators, and captive portal users. No external authentication server is required with this configuration, so all account management is performed on the device or from Panorama.

When configuring Captive Portal, you first create the local account, add it to a User Group and create an Authentication Profile using the new group. You then enable Captive Portal from Device > User Authentication > Captive Portal and select the Authentication Profile. Once this is configured, you can then create a policy from Policies > Captive Portal. For more information, refer to "Captive Portals" on page 285.

### Adding Local Users

▶ *Device > Local User Database > Users*

Use the **Local Users** page to add user information to the local database.

**Table 13. Local User Settings**

| Field | Description |
|---|---|
| Local User Name | Enter a name to identify the user (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Location | Choose a virtual system or choose **Shared** to make the certificate available to all virtual systems. |
| Mode | Use this field to specify the authentication option:<br>• **Password**—Enter and confirm a password for the user.<br>• **Password Hash**—Enter a hashed password string. |
| Enable | Select the check box to activate the user account. |

### Adding Local User Groups

▶ *Device > Local User Database > User Groups*

Use the **Local User Groups** page to add user group information to the local database.

**Table 14. Local User Group Settings**

| Field | Description |
|---|---|
| Local User Group Name | Enter a name to identify the group (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Location | Choose a virtual system or choose **Shared** to make the certificate available to all virtual systems. |
| All Local Users | Click **Add** to select the users you want to add to the group. |

## Configuring RADIUS Server Settings

▶ *Device > Server Profiles > RADIUS*

Use the **RADIUS** page to configure settings for the RADIUS servers that are identified in authentication profiles. Refer to "Authentication Profiles" on page 62.

**Table 15. RADIUS Server Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the server (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Location | Choose a virtual system, or choose **Shared** to make the profile available to all virtual systems. |
| Administrator Use Only | Use this server profile for administrator authentication only. |
| Domain | Enter the RADIUS server domain. The domain setting is used if the user does not specify a domain when logging in. |
| Timeout | Enter an interval after which an authentication request times out (1-30 seconds, default 3 seconds). |
| Retries | Enter the number of automatic retries following a timeout before the request fails (1-5, default 3). |
| Retrieve User Group | Select the check box to use RADIUS VSAs to define the group that has access to the firewall. |
| Servers | Configure information for each server in the preferred order.<br>• **Name**—Enter a name to identify the server.<br>• **IP address**—Enter the server IP address.<br>• **Port**—Enter the server port for authentication requests.<br>• **Secret/Confirm Secret**—Enter and confirm a key to verify and encrypt the connection between the firewall and the RADIUS server. |

## Configuring LDAP Server Settings

▶ *Device > Server Profiles > LDAP*

Use the **LDAP** page to configure settings for the LDAP servers to use for authentication by way of authentication profiles. Refer to "Authentication Profiles" on page 62.

**Table 16. LDAP Server Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Location | Choose a virtual system, or choose **Shared** to make the profile available to all virtual systems. |
| Administrator Use Only | Use this server profile for administrator authentication only. |
| Servers | Specify the host names, IPv4 or IPv6 addresses, and ports of your LDAP servers. |
| Domain | Enter the server domain name. This domain name should be the NetBIOS name of the domain and will be added to the username when authentication is performed. For example, if your domain is paloaltonetworks.com, you only need to enter paloaltonetworks. |
| Type | Choose the server type from the drop-down list. |
| Base | Specify the root context in the directory server to narrow the search for user or group information. |
| Bind DN | Specify the login name (Distinguished Name) for the directory server. |
| Bind Password / Confirm Bind Password | Specify the bind account password. The agent saves the encrypted password in the configuration file. |
| SSL | Select to use secure SSL or Transport Layer Security (TLS) communications between the Palo Alto Networks device and the directory server. |
| Time Limit | Specify the time limit imposed when performing directory searches (1 - 30 seconds, default 30 seconds). |
| Bind Time Limit | Specify the time limit imposed when connecting to the directory server (1 - 30 seconds, default 30 seconds). |
| Retry Interval | Specify the interval after which the system will try to connect to the LDAP server after a previous failed attempt (1-3600 seconds). |

### Configuring Kerberos Settings (Native Active Directory Authentication)

▶ *Device > Server Profiles > Kerberos*

Use the **Kerberos** page to configure Active Directory authentication without requiring customers to start Internet Authentication Service (IAS) for RADIUS support. Configuring a Kerberos server allows users to authenticate natively to a domain controller.

When the Kerberos settings are configured, Kerberos becomes available as an option when defining authentication profiles. Refer to "Authentication Profiles" on page 62.

You can configure the Kerberos settings to recognize a user account in any of the following formats, where domain and realm are specified as part of the Kerberos server configuration:

- domain\username
- username@realm
- username

**Table 17. Kerberos Server Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the server (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Location | Choose a virtual system, or choose **Shared** to make the profile available to all virtual systems. |
| Administrator Use Only | Use this server profile for administrator authentication only. |
| Realm | Specify the hostname portion of the user login name (up to 127 characters)<br>Example: The user account name *user@example.local* has realm *example.local*. |
| Domain | Specify the domain for the user account (up to 63 characters). |
| Servers | For each Kerberos server, click **Add** and specify the following settings:<br>• **Server**—Enter the server IP address.<br>• **Host**—Enter the server FQDN.<br>• **Port**—Enter an optional port number for communication with the server. |

# Authentication Sequence

In some environments, user accounts reside in multiple directories. Guest or other accounts may also be stored in different directories. An authentication sequence is a set of authentication profiles that are applied in order when a user attempts to log in to the firewall. The firewall will always try the local database first, and then each profile in sequence until the user is identified. Access to the firewall is denied only if authentication fails for any of the profiles in the authentication sequence.

For example, after the local database is checked, you can configure an authentication sequence to then try RADIUS, followed by LDAP authentication.

## Setting Up Authentication Sequences

▶ *Device > Authentication Sequence*

Use the **Authentication Sequence** page to configure sets of authentication profiles that are tried in order when a user requests access to the firewall. The user is granted access if authentication is successful using any one of the authentication profiles in the sequence. For more information, refer to "Authentication Profiles" on page 62.

**Table 18. Authentication Sequence Settings**

| Field | Description |
|---|---|
| Profile Name | Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems. |
| Lockout Time | Enter the number of minutes that a user is locked out if the number of failed attempts is reached (0-60 minutes, default 0). 0 means that the lockout is in effect until it is manually unlocked. |
| Failed Attempts | Enter the number of failed login attempts that are allowed before the account is locked out (1-10, default 0). 0 means that there is no limit. |
| Profile List | Choose the authentication profiles to include in the authentication sequence. To change the list order, select an entry and click **Move Up** or **Move Down**. |

# Firewall Logs

▶ *Monitor > Logs*

The firewall provides logs that record configuration changes, system events, security threats, and traffic flows. For each log, you can enable remote logging to a Panorama server, and generate SNMP traps, syslog messages, and email notifications.

The following table describes the logs and logging options.

**Table 19 Log Types and Settings**

| Log | Description |
|---|---|
| Alarms | The alarms log records detailed information on alarms that are generated by the system. The information in this log is also reported in the **Alarms** window. Refer to "Viewing Alarms" on page 85. |
| Configuration | The configuration log records each configuration change, including the date and time, the administrator user name, and whether the change succeeded or failed.<br>All configuration log entries can be sent to Panorama, syslog, and email servers, but they cannot generate SNMP traps.<br>To view expanded details about configuration log entries, move your cursor over the **Before Change** of **After Change** column and click the ellipsis symbol ... . A pop-up window opens to show the full details of the entry. |

**Table 19 Log Types and Settings (Continued)**

| Log | Description |
|---|---|
| Data Filtering | The data filtering log records information on the security policies that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall (refer to "Data Filtering Profiles" on page 223.<br>If you configure a file blocking profile to block specific file types, the file type and file name will appear in the data filtering log, so you can see what was blocked. |
| HIP Match | The HIP match log lists the host information profile (HIP) match requests for GlobalProtect. Refer to "Configuring GlobalProtect" on page 335. |
| System | The system log records each system event, such as HA failures, link status changes, and administrators logging in and out. Each entry includes the date and time, the event severity, and an event description.<br>System log entries can be logged remotely by severity level. For example, you can generate SNMP traps and email notifications for just critical and high-level events. |
| Threat | The threat log records each security alarm generated by the firewall. Each entry includes the date and time, the threat type, such as a virus or spyware/vulnerability filtering violation, the source and destination zones, addresses, and ports, the application name, and the action and severity.<br>Threat log entries can be logged remotely by severity level by defining log forwarding profiles, and then assigning the profiles to security rules (refer to "Security Policies" on page 187). Threats are logged remotely only for the traffic that matches the security rules where the logging profile is assigned.<br>Threat logs are used in generating reports and in the Application Command Center (refer to "Reports and Logs" on page 251). |
| Traffic | The traffic log can record an entry for the start and end of each session. Each entry includes the date and time, the source and destination zones, addresses, and ports, the application name, the security rule applied to the session, the rule action (allow, deny, or drop), the ingress and egress interface, and the number of bytes.<br>Each security rule specifies whether the start and/or end of each session is logged locally for traffic that matches the rule. The log forwarding profile assigned to the rule determines whether the locally logged entries are also logged remotely.<br>Traffic logs are used in generating reports and in the Application Command Center (refer to "Reports and Logs" on page 251). |
| URL Filtering | The URL filtering log records entries for URL filters, which block access to specific web sites and web site categories or generate an alert when a user accesses a proscribed web site.<br>If you are using the PAN-DB and wold like to request a URL category change for a website, open the URL Filtering log details for the desired log entry and click on **Request Categorization Change**. Select the suggested category from the drop-down menu, enter your email address (optional) and comments (optional) and click **Send**.<br>If you are using the BrightCloud DB, click **Request Categorization Change** and you will be directed to the BrightCloud support page where you can submit you request.<br>(refer to "URL Filtering Profiles" on page 217). |

**Table 19 Log Types and Settings (Continued)**

| Log | Description |
|---|---|
| WildFire | The WildFire logs will show the result of files analyzed by WildFire. Some of the fields that will be displayed include: Filename, Source Zone, Destination Zone, Attacker, Victim, and Application.<br>***Note:*** *A WildFire subscription is required for integrated logging. If you do not have a subscription, you can view log information by using the WildFire portal at https://wildfire.paloaltonetworks.com.*<br>(refer to "About WildFire" on page 421 and "Using the WildFire Portal" on page 425) |

# Logging Configuration

You can configure the firewall to send log entries to a Panorama centralized management system, SNMP trap sinks, syslog servers, and email addresses.

The following table describes the remote log destinations.

**Table 20 Remote Log Destinations**

| Destination | Description |
|---|---|
| Panorama | All log entries can be forwarded to a Panorama centralized management system. To specify the address of the Panorama server, refer to "Defining Management Settings" on page 30. |
| SNMP trap | SNMP traps can be generated by severity level for system, threat, and traffic log entries, but not for configuration log entries. To define the SNMP trap destinations, refer to "Configuring SNMP Trap Destinations" on page 75. |
| Syslog | Syslog messages can be generated by severity level for system, threat, and traffic log entries, and for all configuration log entries. To define the syslog destinations, refer to "Configuring Syslog Servers" on page 76. |
| Email | Emails can be generated by severity level for system, threat, and traffic log entries, and for all configuration log entries. To define the email addresses and servers, refer to "You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in Device > Server Profiles > Syslog. Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type." on page 77. |

## Scheduling Log Exports

▶ *Device > Scheduled Log Export*

You can schedule exports of logs and save them to a File Transfer Protocol (FTP) server in CSV format or use Secure Copy (SCP) to securely transfer data between the device and a remote host. Log profiles contain the schedule and FTP server information. For example, a profile may specify that the previous day's logs are collected each day at 3AM and stored on a particular FTP server.

When you click **OK** after creating a new entry, the new profile is added to the **Scheduled Log Export** page. You must commit the change for the export to take place. If you are using SCP, you need to click the **Test SCP server connection** button to test connectivity between the firewall and the SCP server and you must verify and accept the host key of the SCP server.

**Table 21. Scheduled Log Export Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.<br>You cannot change the name after the profile is created. |
| Description | Enter an optional description (up to 255 characters). |
| Enabled | Select the check box to enable the scheduling of log exports. |
| Log Type | Select the type of log (traffic, threat, url, data, or hipmatch). Default is traffic. |
| Scheduled export start time (daily) | Enter the time of day (hh:mm) to start the export, using a 24-hour clock (00:00 - 23:59). |
| Protocol | Select the protocol to use to export logs from the firewall to a remote host. You can use SCP to export logs securely, or you can use FTP, which is not a secure protocol. |
| Hostname | Enter the host name or IP address of the FTP server that will be used for the export. |
| Port | Enter the port number that the FTP server will use. Default is 21. |
| Path | Specify the path located on the FTP server that will be used to store the exported information. |
| Enable FTP Passive Mode | Select the check box to use passive mode for the export. By default, this option is selected. |
| Username | Enter the user name for access to the FTP server. Default is anonymous. |
| Password | Enter the password for access to the FTP server. A password is not required if the user is "anonymous." |

## Defining Configuration Log Settings

▶ *Device > Log Settings > Config*

The configuration log settings specify the configuration log entries that are logged remotely with Panorama, and sent as syslog messages and/or email notifications.

**Table 22. Configuration Log Settings**

| Field | Description |
|---|---|
| Panorama | Select the check box to enable sending configuration log entries to the Panorama centralized management system. |
| SNMP Trap | To generate SNMP traps for configuration log entries, select trap name. To specify new SNMP trap destinations, refer to "Configuring SNMP Trap Destinations" on page 75. |
| Email | To generate email notifications for configuration log entries, select an email profile from the drop-down menu. To specify a new email profile, refer to "Configuring Email Notification Settings" on page 83. |
| Syslog | To generate syslog messages for configuration log entries, select the name of the syslog server. To specify new syslog servers, refer to "Configuring Syslog Servers" on page 76. |

## Defining System Log Settings

▶ *Device > Log Settings > System*

The system log settings specify the severity levels of the system log entries that are logged remotely with Panorama and sent as SNMP traps, syslog messages, and/or email notifications. The system logs show system events such as HA failures, link status changes, and administrators logging in and out.

**Table 23. System Log Settings**

| Field | Description |
|---|---|
| Panorama | Select the check box for each severity level of the system log entries to be sent to the Panorama centralized management system. To specify the Panorama server address, refer to "Defining Management Settings" on page 30.<br>The severity levels are:<br>• **Critical**—Hardware failures, including HA failover, and link failures.<br>• **High**—Serious issues, including dropped connections with external devices, such as syslog and RADIUS servers.<br>• **Medium**—Mid-level notifications, such as antivirus package upgrades.<br>• **Low**—Minor severity notifications, such as user password changes.<br>• **Informational**—Log in/log off, administrator name or password change, any configuration change, and all other events not covered by the other severity levels. |

**Table 23. System Log Settings (Continued)**

| Field | Description |
|---|---|
| SNMP Trap<br>Email<br>Syslog | Under each severity level, select the SNMP, syslog, and/or email settings that specify additional destinations where the system log entries are sent. To define new destinations, refer to:<br>• "Configuring SNMP Trap Destinations" on page 75.<br>• "Configuring Syslog Servers" on page 76.<br>• "Configuring Email Notification Settings" on page 83. |

## Defining HIP Match Log Settings

▶ *Device > Log Settings > HIP Match*

The Host Information Profile (HIP) match log settings are used to provide information on security policies that apply to GlobalProtect clients. For more information, refer to "Overview" on page 335.

**Table 24. HIP Match Log Settings**

| Field | Description |
|---|---|
| Panorama | Select the check box to enable sending configuration log entries to the Panorama centralized management system. |
| SNMP Trap | To generate SNMP traps for HIP match log entries, select the name of the trap destination. To specify new SNMP trap destinations, refer to "Configuring SNMP Trap Destinations" on page 75. |
| Email | To generate email notifications for configuration log entries, select the name of the email settings that specify the appropriate email addresses. To specify new email settings, refer to "You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in Device > Server Profiles > Syslog. Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type." on page 77. |
| Syslog | To generate syslog messages for configuration log entries, select the name of the syslog server. To specify new syslog servers, refer to "Configuring Syslog Servers" on page 76. |

## Defining Alarm Log Settings

▶ *Device > Log Settings > Alarms*

Use the **Alarms** page to configure notifications when a security rule (or group of rules) has been hit repeatedly in a set period of time.

**Table 25. Alarm Log Settings**

| Field | Description |
|---|---|
| Enable Alarms | Enable alarms based on the events listed on this page. |
| Enable CLI Alarm Notifications | Enable CLI alarm notifications whenever alarms occur. |

**Table 25. Alarm Log Settings (Continued)**

| Field | Description |
|---|---|
| Enable Web Alarm Notifications | Open a window to display alarms on user sessions, including when they occur and when they are acknowledged. |
| Enable Audible Alarms | The firewall will continuously play an audible tone when unacknowledged alarms exist in the web interface or CLI. |
| Encryption/Decryption Failure Threshold | Specify the number of encryption/decryption failures after which an alarm is generated. |
| Log DB Alarm Threshold (% Full) | Generate an alarm when a log database reaches the indicated percentage of the maximum size. |
| Security Policy Limits | An alarm is generated if a particular IP address or port hits a deny rule the number of times specified in the **Security Violations Threshold** setting within the period (seconds) specified in the **Security Violations Time Period** setting. |
| Security Policy Group Limits | An alarm is generated if the collection of rules reaches the number of rule limit violations specified in the **Violations Threshold** field during the period specified in the **Violations Time Period** field. Violations are counted when a session matches an explicit deny policy.<br>Use **Security Policy Tags** to specify the tags for which the rule limit thresholds will generate alarms. These tags become available to be specified when defining security policies. |
| Selective Audit | ***Note:*** *These settings appear on the* ***Alarms*** *page only in Common Criteria mode.*<br>Specify the following settings:<br>• **CC Specific Logging**—Enables verbose logging required for Common Criteria (CC) compliance.<br>• **Login Success Logging**—Logs the success of administrator logins to the firewall.<br>• **Login Failure Logging**—Logs the failure of administrator logins to the firewall.<br>• **Suppressed Administrators**—Does not generate logs for changes that the listed administrators make to the firewall configuration. |

## Managing Log Settings

▶ ***Device > Log Settings > Manage Logs***

Click the links on this page to clear the indicated logs.

# Configuring SNMP Trap Destinations

▶ *Device > Server Profiles > SNMP Trap*

To generate SNMP traps for system, traffic, or threat logs, you must specify one or more SNMP trap destinations. After you define the trap destinations, you can use them for system log entries (refer to "Defining System Log Settings" on page 72).

**Table 26. SNMP Trap Destination Settings**

| Field | Description |
|---|---|
| Name | Enter a name for the SNMP profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems. |
| Version | Choose the SNMP version (V2c or V3). |
| V2c settings | If you choose V2c, configure the following settings:<br>• **Server**—Specify a name for the SNMP trap destination name (up to 31 characters).<br>• **Manager**—Specify the IP address of the trap destination.<br>• **Community**—Specify the community string required to send traps to the specified destination (default public). |
| V3 settings | If you choose V3, configure the following settings:<br>• **Server**—Specify the SNMP trap destination name (up to 31 characters).<br>• **Manager**—Specify the IP address of the trap destination.<br>• **User**—Specify the SNMP user.<br>• **EngineID**—Specify the engine ID of the firewall. The input is a string in hexadecimal representation. The engine ID is any number between 5 to 64 bytes. When represented as a hexadecimal string this is between 10 to 128 characters (2 characters for each byte) with two additional characters for 0x that you need to use as a prefix in the input string.<br>Each firewall has a unique engine ID, which you can get by using a MIB browser to run a GET for OID 1.3.6.1.6.3.10.2.1.1.0.<br>• **Auth Password**—Specify the user's authentication password (minimum 8 characters, maximum of 256 characters, and no character restrictions). All characters allowed). Only Secure Hash Algorithm (SHA) is supported.<br>• **Priv Password**—Specify the user's encryption password (minimum 8 characters, maximum of 256 characters, and no character restrictions). Only Advanced Encryption Standard (AES) is supported. |

***Note:*** *Do not delete a destination that is used in any system log settings or logging profile.*

### SNMP MIBs

The firewall supports the following SNMP MIBs:

- "RFC 1213: MIB-II - Support for The System Group, The Interfaces Group.
- "RFC 2863: IF-MIB - The Interfaces Group MIB
- "RFC 2790: HOST-RESOURCES-MIB - Support for hrDeviceTable and hrProcessorTable.
- "RFC 3433: ENTITY-SENSOR-MIB - Support for entPhySensorTable.
- PAN-PRODUCT-MIB
- PAN-COMMON-MIB
- PAN-TRAPS-MIB
- PAN-LC-MIB

The full set of Enterprise MIBs is available in the Technical Documentation section on the Palo Alto Networks site at *https://live.paloaltonetworks.com/community/documentation*.

# Configuring Syslog Servers

▶ ***Device > Server Profiles > Syslog***

To generate syslog messages for system, configuration, traffic, threat, or HIP match logs, you must specify one or more syslog servers. After you define the syslog servers, you can use them for system and configuration log entries (refer to "Defining System Log Settings" on page 72).

**Table 27. New Syslog Server**

| Field | Description |
|---|---|
| Name | Enter a name for the syslog profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems. |
| **Servers Tab** | |
| Name | Click **Add** and enter a name for the syslog server (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Server | Enter the IP address of the syslog server. |
| Port | Enter the port number of the syslog server (the standard port is 514). |
| Facility | Choose a level from the drop-down list. |

**Table 27. New Syslog Server (Continued)**

| Field | Description |
|---|---|
| **Custom Log Format Tab** | |
| Log Type | Click the log type to open a dialog box that allows you to specify a custom log format. In the dialog box, click a field to add it to the Log Format area. Other text strings can be edited directly in the Log Format area. Click **OK** to save the settings.<br>For details on the fields that can be used for custom logs, refer to "Custom Syslog Field Descriptions" on page 77. |
| Escaping | Specify escape sequences. Use the **Escaped characters** box to list all the characters to be escaped without spaces. |

***Note:*** *You cannot delete a server that is used in any system or configuration log settings or logging profiles.*

## Custom Syslog Field Descriptions

You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in **Device > Server Profiles > Syslog.** Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type.

**Table 28 Config Fields**

| Field | Meaning |
|---|---|
| actionflags | A bit field indicating if the log was forwarded to Panorama. Available in PAN-OS 4.0.0 and above. |
| admin | User name of the Administrator performing the configuration. |
| after-change-detail | Details of the configuration after a change is made. |
| before-change-detail | Details of the configuration before a change is made. |
| formtted-receive_time | Time the log was received at the management plane, shown in CEF compliant time format. |
| cef-formatted-time_generated | Time the log was generated, shown in CEF compliant time format. |
| client | Client used by the Admin; Values are Web and CLI. |
| cmd | Command performed by the Admin; Values are add, clone, commit, delete, edit, move, rename, set, validate. |
| host | Host name or IP address of the client machine |
| path | The path of the configuration command issued. Up to 512 bytes in length. |
| receive_time | Time the log was received at the management plane. |
| result | Result of the configuration action. Values are Submitted, Succeeded, Failed, and Unauthorized. |

**Table 28 Config Fields**

| Field | Meaning |
|---|---|
| seqno | A 64bit log entry identifier incriminated sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above. |
| serial | Serial number of the device that generated the log |
| subtype | Subtype of the Config log; Unused. |
| time_generated | Time the log was received on the data plane. |
| type | Specifies type of log; Values are traffic, threat, config, system and hip-match. |
| vsys | Virtual System associated with the configuration log. |

**Table 29 System Fields**

| Field | Meaning |
|---|---|
| actionflags | A bit field indicating if the log was forwarded to Panorama. Available in PAN-OS 4.0.0 and above. |
| cef-formatted-receive_time | Time the log was received at the management plane, shown in CEF compliant time format. |
| cef-formatted-time_generated | Time the log was generated, shown in CEF compliant time format. |
| eventid | String showing the name of the event |
| fmt | Detailed description of the event. Length is up to 512 bytes. |
| module | This field is valid only when the value of the Subtype field is general; It provides additional information about the sub-system generating the log. Values are general, management, auth, ha, upgrade, chassis. |
| number-of-severity | Severity level as an integer - informational-1, low-2, medium-3, high-4, critical-5. |
| object | Name of the object associated with the system log. |
| opaque | Detailed description of the event. Length is up to 512 bytes. |
| receive_time | Time the log was received at the management plane |
| seqno | A 64bit log entry identifier that increments sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above. |
| serial | Serial number of the device that generated the log |
| severity | Severity associated with the event; Values are informational, low, medium, high, critical |
| subtype | Subtype of the system log. Refers to the system daemon generating the log; Values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn. |
| time_generated | Time the log was received on the data plane. |

**Table 29 System Fields**

| Field | Meaning |
|---|---|
| type | Specifies type of log; Values are traffic, threat, config, system and hip-match. |
| vsys | Virtual System associated with the system event |

**Table 30 Threat Fields**

| Field | Meaning |
|---|---|
| action | Action taken for the session; Values are alert, allow, deny, drop, drop-all-packets, reset-client, reset-server, reset-both, block-url. See Action Field table below for meaning of each value. |
| actionflags | A bit field indicating if the log was forwarded to Panorama. Available in PAN-OS 4.0.0 and above. |
| app | Application associated with the session. |
| category | For URL Subtype, it is the URL category; for WildFire subtype, it is the verdict on the file and is either “malicious” or “benign”; for other subtypes, the value is “any”. |
| cef-formatted-receive_time | Time the log was received at the management plane, shown in CEF compliant time format. |
| cef-formatted-time_generated | Time the log was generated, shown in CEF compliant time format. |
| contenttype | Content type of the HTTP response data. Maximum length 32 bytes. Applicable only when Subtype is URL. Available in PAN-OS 4.0.0 and above. |
| direction | Indicates the direction of the attack, 'client-to-server' or 'server-to-client'. |
| dport | Destination port utilized by the session. |
| dst | Original session destination IP address. |
| dstloc | Destination country or Internal region for private addresses. Maximum length is 32 bytes. Available in PAN-OS 4.0.0 and above. |
| dstuser | User name of the user to which the session was destined. |
| flags | 32 bit field that provides details on the session; See Flags Field table for meaning of each value. |
| from | Zone the session was sourced from. |
| inbound_if | Interface that the session was sourced form. |
| logset | Log Forwarding Profile that was applied to the session. |
| misc | The actual URI when the subtype is URL; File name or file type when the subtype is file; and File name when the subtype is virus; File name when the subtype is wildfire. Length is 63 characters in PAN-OS versions before 4.0. From version 4.0, it is variable length with a maximum of 1023 characters. |
| natdport | Post-NAT destination port. |

**Table 30 Threat Fields**

| Field | Meaning |
|---|---|
| natdst | If Destination NAT performed, the post-NAT Destination IP address. |
| natsport | Post-NAT source port. |
| natsrc | If Source NAT performed, the post-NAT Source IP address. |
| number-of-severity | Severity level as an integer - informational-1, low-2, medium-3, high-4, critical-5. |
| outbound_if | Interface that the session was destined to. |
| proto | IP protocol associated with the session. |
| receive_time | Time the log was received at the management plane. |
| repeatcnt | Number of logs with same Source IP, Destination IP, and Threat ID seen within 5 seconds; Applies to all Subtypes except URL. |
| rule | Name of the rule that the session matched. |
| seqno | A 64bit log entry identifier that increments sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above. |
| serial | Serial number of the device that generated the log. |
| sessionid | An internal numerical identifier applied to each session. |
| severity | Severity associated with the threat; Values are informational, low, medium, high, critical. |
| sport | Source port utilized by the session. |
| src | Original session source IP address. |
| srcloc | Source country or Internal region for private addresses. Maximum length is 32 bytes. Available in PAN-OS 4.0.0 and above. |
| srcuser | User name of the user that initiated the session. |
| subtype | Subtype of threat log; Values are URL, virus, spyware, vulnerability, file, scan, flood, data, and wildfire. |
| threatid | Palo Alto Networks identifier for the threat. It is a description string followed by a numerical identifier in parenthesis for some Subtypes. The numerical identifier is a 64-bit number from PAN-OS 5.0 and later. |
| time_generated | Time the log was generated on the data plane. |
| time_received | Time the log was received on the data plane. |
| to | Zone the session was destined to. |
| type | Specifies type of log; Values are traffic, threat, config, system and hip-match. |
| vsys | Virtual System associated with the session. |
| wildfire | Logs generated by WildFire. |

**Table 31 Traffic Fields**

| Field | Meaning |
|---|---|
| action | Action taken for the session; Values are allow or deny. See Action Field table. |
| actionflags | A bit field indicating if the log was forwarded to Panorama. Available from PAN-OS 4.0.0. |
| app | Application associated with the session. |
| bytes | Number of total bytes (transmit and receive) for the session. |
| bytes_received | Number of bytes in the server-to-client direction of the session. Available from PAN-OS 4.1.0 on all models except the PA-4000 series. |
| bytes_sent | Number of bytes in the client-to-server direction of the session. Available from PAN-OS 4.1.0 on all models except the PA-4000 series. |
| category | URL category associated with the session (if applicable). |
| cef-formatted-receive_time | Time the log was received at the management plane, shown in CEF compliant time format. |
| cef-formatted-time_generated | Time the log was generated, shown in CEF compliant time format. |
| dport | Destination port utilized by the session. |
| dst | Original session destination IP address. |
| dstloc | Destination country or Internal region for private addresses. Maximum length is 32 bytes. Available in PAN-OS 4.0.0 and above. |
| dstuser | User name of the user to which the session was destined. |
| elapsed | Elapsed time of the session. |
| flags | 32 bit field that provides details on session; See Flags Field table for meaning of each value. This field can be decoded by AND-ing the values with the logged value. |
| from | Zone the session was sourced from. |
| inbound_if | Interface that the session was sourced form. |
| logset | Log Forwarding Profile that was applied to the session. |
| natdport | Post-NAT destination port. |
| natdst | If Destination NAT performed, the post-NAT Destination IP address. |
| natsport | Post-NAT source port. |
| natsrc | If Source NAT performed, the post-NAT Source IP address. |
| outbound_if | Interface that the session was destined to. |
| packets | Number of total packets (transmit and receive) for the session. |
| pkts_received | Number of server-to-client packets for the session. Available from PAN-OS 4.1.0 on all models except the PA-4000 series. |
| pkts_sent | Number of client-to-server packets for the session. Available from PAN-OS 4.1.0 on all models except the PA-4000 series. |

**Table 31 Traffic Fields**

| Field | Meaning |
|---|---|
| proto | IP protocol associated with the session. |
| receive_time | Time the log was received at the management plane. |
| repeatcnt | Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; Used for ICMP only. |
| rule | Name of the rule that the session matched. |
| seqno | A 64bit log entry identifier incremented sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above. |
| serial | Serial number of the device that generated the log. |
| sessionid | An internal numerical identifier applied to each session. |
| sport | Source port utilized by the session. |
| src | Original session source IP address. |
| srcloc | Source country or Internal region for private addresses. Maximum length is 32 bytes. Available in PAN-OS 4.0.0 and above. |
| srcuser | User name of the user that initiated the session. |
| start | Time of session start. |
| subtype | Subtype of traffic log; Values are start, end, drop, and deny. See Subtype Field table for meaning of each value. |
| time_generated | Time the log was generated on the data plane. |
| time_received | Time the log was received on the data plane. |
| to | Zone the session was destined to. |
| type | Specifies type of log; Values are traffic, threat, config, system and hip-match. |
| vsys | Virtual System associated with the session. |

**Table 32 HIP Match Fields**

| Field | Meaning |
|---|---|
| actionflags | A bit field indicating if the log was forwarded to Panorama. Available in PAN-OS 4.0.0 and above. |
| cef-formatted-receive_time | Time the log was received at the management plane, shown in CEF compliant time format. |
| cef-formatted-time_generated | Time the log was generated, shown in CEF compliant time format. |
| machinename | Name of the Users machine. |
| matchname | Name of the HIP Object or Profile. |
| matchtype | Specifies whether the HIP field represents a HIP Object or a HIP Profile. |
| receive_time | Time the log was received at the management plane. |

**Table 32 HIP Match Fields**

| Field | Meaning |
|---|---|
| repeatcnt | Number of times the HIP profile matched. |
| seqno | A 64bit log entry identifier incremented sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above. |
| serial | Serial number of the device that generated the log. |
| src | IP address of the source user. |
| srcuser | User name of the Source user. |
| subtype | Subtype of hip-match log; Unused. |
| time_generated | Time the log was generated on the data plane. |
| type | Specifies type of log; Values are traffic, threat, config, system and hip-match. |
| vsys | Virtual System associated with the HIP Match log. |

# Configuring Email Notification Settings

▶ *Device > Server Profiles > Email*

To generate email messages for logs, you must configure an email profile. After you define the email settings, you can enable email notification for system and configuration log entries (refer to "Defining System Log Settings" on page 72). For information on scheduling email report delivery, refer to "Scheduling Reports for Email Delivery" on page 273.

**Table 33. Email Notification Settings**

| Field | Description |
|---|---|
| Name | Enter a name for the email settings (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| **Servers Tab** | |
| Server | Enter a name to identify the server (1-31 characters). This field is just a label and does not have to be the host name of an existing SMTP server. |
| Display Name | Enter the name shown in the From field of the email. |
| From | Enter the From email address, such as "security_alert@company.com". |
| To | Enter the email address of the recipient. |
| Additional Recipient | Optionally, enter the email address of another recipient. You can only add one additional recipient. To add multiple recipients, add the email address for a distribution list. |
| Gateway | Enter the IP address or host name of the Simple Mail Transport Protocol (SMTP) server used to send the email. |

**Table 33. Email Notification Settings (Continued)**

| Field | Description |
|---|---|
| **Custom Log Format Tab** | |
| Log Type | Click the log type to open a dialog box that allows you to specify a custom log format. In the dialog box, click a field to add it to the Log Format area. Click **OK** to save the settings. |
| Escaping | Include escaped characters and specify the escape character or characters. |

***Note:*** *You cannot delete an email setting that is used in any system or configuration log settings or logging profiles.*

# Viewing Alarms

You can view the current list of alarms at any time by clicking the **Alarms** icon in the lower right corner of the web interface when the Alarm option is configured. This opens a window that lists the unacknowledged and acknowledged alarms in the current alarms log. To acknowledge alarms, select their check boxes and click **Acknowledge**. This action moves the alarms to the Acknowledged Alarms list. The alarms window also includes paging, column sort, and refresh controls.

The Alarms button is visible only when the **Enable Alarms** check box is selected on the **Device > Log Settings > Alarms > Alarm Settings** page.

# Configuring Netflow Settings

▶ *Device > Server Profiles > Netflow*

The firewall can generate and export NetFlow Version 9 records with unidirectional IP traffic flow information to an outside collector. NetFlow export can be enabled on any ingress interface in the system. Separate template records are defined for IPv4, IPv4 with NAT, and IPv6 traffic, and PAN-OS specific fields for App-ID and User-ID can be optionally exported. This feature is available on all platforms, except for the PA-4000 Series models.

The firewall supports the standard NetFlow templates and selects the correct one based on the data to be exported.

To configure NetFlow data exports, define a NetFlow server profile, which specifies the frequency of the export along with the NetFlow servers that will receive the exported data.

Then when you assign the profile to an existing firewall interface, all traffic flowing over that interface is exported to the specified servers. All interface types support assignment of a NetFlow profile. Refer to "Firewall Interfaces" on page 127 for information on assigning a NetFlow profile to an interface.

**Table 34. Netflow Settings**

| Field | Description |
|---|---|
| Name | Enter a name for the Netflow settings (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Template Refresh Rate | Specify the number of minutes or number of packets after which the Netflow template is refreshed (minutes range 1-3600, default 30 min; packets range 1-600, default 20). |
| Active Timeout | Specify the frequency at which data records are exported for each session (minutes). |
| Export PAN-OS Specific Field Types | Export PAN-OS specific fields such as App-ID and User-ID in Netflow records. |
| **Servers** | |
| Name | Specify a name to identify the server (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |

**Table 34. Netflow Settings (Continued)**

| Field | Description |
|---|---|
| Server | Specify the host name or IP address of the server. You can add a maximum of two servers per profile. |
| Port | Specify the port number for server access (default 2055). |

# Importing, Exporting and Generating Security Certificates

## Certificates

▶ *Device > Certificate Management > Certificates*

The **Certificates** page allows you to generate the following security certificates:

- **Forward Trust**—This certificate is presented to clients during decryption when the server to which they are connecting is signed by a CA in the firewall's trusted CA list. If a self-signed certificate is used for forward proxy decryption, you must click the certificate name in the **Certificates** page and select the **Forward Trust Certificate** check box.
- **Forward Untrust**—This certificate is presented to clients during decryption when the server to which they are connecting is signed by a CA that is not in the firewall's trusted CA list.
- **Trusted Root CA**—The certificate is marked as a trusted CA for forward decryption purposes.

  When the firewall decrypts traffic, it checks the upstream certificate to see if it is issued by a trusted CA. If not, it uses a special untrusted CA certificate to sign the decryption certificate. In this case, the user sees the usual certificate error page when accessing the firewall and must dismiss the login warning.

  The firewall has a large list of existing trusted CAs. The trusted root CA certificate is for additional CAs that are trusted for your enterprise but are not part of the pre-installed trusted list.
- **SSL Exclude**—This certificate excludes connections if they are encountered during SSL forward proxy decryption.
- **Certificate for Secure Web GUI**—This certificate authenticates users for access to the firewall web interface. If this check box is selected for a certificate, the firewall will use this certificate for all future web-based management sessions following the next commit operation.

Perform any of the following functions on the **Certificates** page:

- To revoke a certificate:

  a. Select the certificate that you want to revoke.

b. Click **Revoke** and the certificate will be instantly set to the revoked status. No commit is required.

   You can also click an existing certificate and click the **Revoke** icon.

- To renew a certificate:

  a. Select the certificate that you want to renew.

  b. Click **Renew** and then set the number of days to extend the certificate and click **OK**. A new certificate will be generated with the same attributes, but with a new serial number. The old certificate is then replaced with the new certificate.

- To import a web interface, trusted CA, or SSL forward proxy certificate:

  a. Click **Import**.

  b. Enter a name to identify the certificate.

  c. Select the certificate file. If importing a PKCS #12 certificate and private key, this will be the single file holding both objects. If using PEM, this will be the public certificate only.

  d. Click the **Import Private Key** check box to load the private key and enter the passphrase twice. If using the PKCS #12, the key file was selected above. If using PEM, browse to the encrypted private key file (generally named *.key).

  e. Select the virtual system to which you want to import the certificate from the drop-down list.

- To export a certificate:

  a. Select the certificate you want to export.

  b. Click **Export**.

  c. Choose the file format you would like the exported certificate to use (.pfx for PKCS#12 or .pem).

  d. Select the **Export Private Key** check box and enter a passphrase twice to export the private key in addition to the certificate.

- Click **Save** and choose a location to copy the file to your local computer. To generate a certificate:

  a. Click **Generate** to open the Generate Certificate window and specify the information described in the following table.

  b. After generating the certificate, click the certificate link and specify the certificate type (Forward Trust, Forward Untrust, Trusted Root CA, SSL Exclude, or Certificate for Secure Web GUI).

***Note:*** *If you are using Panorama, you also have the option of generating a self-signed certificate for the Panorama server. Refer to "Central Device Management Using Panorama" on page 381 for information on Panorama.*

- To import keys for high availability (HA), click **Import HA Key** and browse to specify the key file for import. To export keys for HA, click **Export HA Key** and specify a location to save the file. The HA keys must be swapped across the two firewalls. In other words, the Key from firewall 1 must be exported and then imported to firewall 2 and vice versa.

**Table 35. Settings to Generate a Certificate**

| Field | Description |
|---|---|
| Certificate Name | Enter a name (up to 31 characters) to identify the certificate. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required. |
| Common Name | Enter the IP address or FQDN that will appear on the certificate. |
| Location | Choose a virtual system or choose **Shared** to make the certificate available to all virtual systems. |
| Signed By | Choose from a list of the CA certificates that were generated on the firewall. The selected certificate can be used to sign the certificate that is being created. |
| Certificate Authority | Mark this certificate as a CA so that it can be used to sign other certificates on the firewall. |
| OCSP Responder | Select an OSCP responder profile from the drop-down list. The profile is configured in Device > Certificate Management > OCSP Responder. When generating a certificate and entering an OCSP Responder, a look up will be performed for the host name of the IP address to generate a OCSP Responder URL, which will then appear in this drop-down. |
| Number of Bits | Choose the key length for the certificate. |
| Digest | Choose the digest algorithm for the certificate. |
| Expiration (days) | Specify the number of days that the certificate will be valid. The default is 365 days.<br>If you specify a **Validity Period** in a GlobalProtect Portal Satellite configuration, that value will override the value entered in this field. |
| Country<br>State<br>Locality<br>Organization<br>Department<br>Email | Optionally specify additional information to identify the certificate.<br>To view a list of country code definitions, click the **ISO 6366 Country Codes** link. |

# Default Trusted Certificate Authorities

▶ ***Device > Certificate Management > Certificates > Default Trusted Certificate Authorities***

Use this page to control the certificate authorities (CAs) that the firewall will trust. You can disable and enable CAs as needed.

**Table 36 Trusted Certificate Authorities Settings**

| Field | Description |
|---|---|
| Enable | If you have disabled a CA and want to enable it, click the check box next to the CA and then click **Enable**. |
| Disable | Click the check box next to the CA that you want to disable, then click **Disable**. This may be desired if you only want to trust certain CAs, or remove all of them to only trust your local CA. |
| Export | Click the check box next to the CA, then click **Export** to export the CA certificate. You can do this to import into another system, or if you want to view the certificate offline. |

# Certificate Profile

▶ *Device > Certificate Management > Certificate Profile*

You can create certificate profiles and then attach a profile to an administrator login on the **Setup** page or to an SSL-VPN login for use in authentication or with captive portals. You can also create certificate profiles to be used by GlobalProtect Gateways for authentication. Refer to "Defining Management Settings" on page 30 and "Captive Portals" on page 285.

**Table 37. Certificate Profile Settings**

| Page Type | Description |
|---|---|
| Name | Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Location | If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems. |
| Username Field | Choose a user name option from the drop-down list. |
| Domain | Enter the domain for the profile. |
| CA Certificates | Choose a CA certificate from the drop-down list, specify the default OCSP URL, select an option to verify the CA certificate, and click **Add**. Repeat to add additional certificates.<br>The **OCSP Verify CA Certificate** drop down is used if you have a separate OCSP responder that can verify certificates. |
| Use CRL | Select the check box to use a Certificate Revocation List (CRL). |
| Use OCSP | Select the check box to use Online Certificate Status Protocol (OCSP) server. OCSP takes precedence over CRL. |
| CRL Receive Timeout | Specify an interval after which CRL requests time out (1 - 60 secs). |
| OCSP Receive Timeout | Specify an interval after which OCSP requests time out (1 - 60 secs). |
| Certificate Status Timeout | Specify an interval after which requests for certificate status time out (1 - 60 secs). |

**Table 37. Certificate Profile Settings (Continued)**

| Page Type | Description |
|---|---|
| Block session if certificate status is unknown | Select the check box to block a sessions if the certificate status is unknown. |
| Block sessions if certificate status cannot be retrieved within timeout | Select the check box to block a session if the certificate status cannot be retrieved within the timeout interval. |

# OCSP Responder

▶ *Device > Certificate Management > OCSP Responder*

Use the **OCSP Responder** (Online Certificate Status Protocol Responder) page to define a server that will be used to verify the revocation status of certificates issues by the PAN-OS device. When generating new certificates, you can specify the OCSP Responder that will be used.

To enable OCSP, go to **Device > Setup > Sessions** and under **Sessions Features** click **Decryption Certificate Revocation Settings**.

**Table 38 OCSP Responder Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the OCSP responder server (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Host Name | Enter the host name of the OCSP responder server that will be used to check certificate revocation status for your devices. |

# Encrypting Private Keys and Passwords on the Firewall

▶ *Device > Master Key and Diagnostics*

Use the **Master Key and Diagnostics** page to specify a master key to encrypt private keys on the firewall. Private keys are stored in encrypted form by default even if a new master key is not specified. If the firewall is in Common Criteria mode, several crypto diagnostics capabilities are available. These diagnostics allow you to run scheduled crypto self-tests and on-demand self-tests.

To create a new master key, enter a 16 character string in the **New Master Key** field and then confirm the key. Enter a lifetime and reminder value and then click **OK**. You will need to update the master key before the expiration. For information on updating master keys, refer to "Updating Master Keys" on page 91.


- "Master Key and Diagnostic Settings" on page 91
- "Updating Master Keys" on page 91

## Master Key and Diagnostic Settings

**Table 39. Master Key and Diagnostics Settings**

| Field | Description |
|---|---|
| Current Master Key | Specify the key that is currently used to encrypt all of the private keys and passwords on the firewall. |
| New Master Key<br>Confirm Master Key | To change the master key, enter and confirm a new key. |
| Life Time | Specify the number of days and hours after which the master key expires (range 1-730 days). |
| Time for Reminder | Specify the number of days and hours before expiration when the user is notified of the impending expiration (range 1-365 days). |
| Common Criteria | In Common Criteria mode, additional buttons are available to run a cryptographic algorithm self-test and software integrity self-test. A scheduler is also included to specify the times at which the two self-tests will run. |

## Updating Master Keys

When creating a new master keys an expiration period will be defined, so it's important to update the keys before the expiration period is reached. When using a master key in a high availability (HA) configuration, it is also important that the same master key is used on both devices to ensure that private keys and certificates are encrypted with the same master key. If the master keys are different, HA configuration synchronization will not work properly. The following section describes how to update the master key if they working properly and instructions are provided if they have already expired or if the master keys on each device in the HA pair are out of sync.
As of PAN-OS 5.0, if the master keys in an HA pair do not match, a critical system log will be generated.

Reasons for updating the master key:

- You want to change the default master key or change a master key that you created.
- Master keys in an HA configuration are out of sync.
- Master key is going to expire soon.

HA Master key update (keys are in sync and have not reached the expiration date):

1. If you are updating the master keys and they are not expired or out of sync, before you update the keys, you need to commit the configuration and ensure that no pending configuration updates exist on both devices in the HA pair.
   You can view the commit status by clicking the **Tasks** link on the bottom right of the web interface on each device. From the CLI, run `show jobs all` to see all running jobs or `show jobs pending` to view pending jobs. To check that no pending updates are present, from configure mode, run `check pending-changes` and you should see **No**.
2. Disable configuration synchronization by navigating to **Device > High Availability** on both devices and from the **General** tab deselect the **Enable Config Sync** check box. Commit the configuration on both devices.
3. Update the master key on device A using 16 characters and commit the configuration.

4. Update the master key on device B with the same master key and commit the configuration.
5. The master keys should now be in sync. Check the logs to make sure no critical system logs related to the master key are present.
6. Enable configuration synchronize again by selecting the **Enable Config Sync** check box and then commit the configuration.

HA master key update (keys out of sync or expired):

1. If the master keys are out of sync or expired, you will see critical errors in the system log. If this occurs, you should immediately disable HA config sync on both devices in the HA pair. Navigate to **Device > High Availability** on both devices and from the **General** tab deselect the **Enable Config Sync** check box. Commit the configuration on both devices.
2. Ensure that no pending configuration updates exist on both devices in the HA pair. If there are pending changes, commit the configuration on both devices and wait until all configuration updates are complete.
   You can view the commit status by clicking the **Tasks** link on the bottom right of the web interface on each device. From the CLI, run `show jobs all` to see all running jobs or `show jobs pending` to view pending jobs. To check that no pending updates are present, from configure mode, run `check pending-changes` and you should see **No**.
3. Update the master key on device A using 16 characters and commit the configuration.
4. Update the master key on device B with the same master key and commit the configuration.
5. Enable configuration synchronization on both devices and commit the configuration.
6. The master keys should now be in sync. Check the logs to make sure no critical system logs related to the master key are present.
7. Enable configuration synchronize again by selecting the **Enable Config Sync** check box and then commit the configuration. If problems persist, please contact the Technical Support department.

# High Availability

PAN-OS supports active/passive and active/active high availability (HA).

***Note:*** *In an HA pair, both firewalls must be the same model and have the same licenses. If state synchronization is enabled, existing sessions continue after a switchover; however, threat prevention functions do not continue. Threat protection will apply to new sessions.*

## Active/Passive HA

In the active/passive configuration, two devices form an HA group to provide redundancy. The two firewalls mirror each other in configuration. If the active firewall fails for any reason, the passive firewall becomes active automatically with no loss of service. A failover can also occur if selected Ethernet links fail or if the active firewall cannot reach one or more of the specified destinations. From a traffic processing perspective, at most one device receives packets at any one time.

The following rules apply to HA operation and failover:

- The active firewall continuously synchronizes its configuration and session information with the passive firewall over the HA interfaces.
- If the active firewall fails, then the passive firewall detects the loss of heartbeats and automatically becomes active.
- If one HA interface fails, synchronization continues over the remaining interface. If the state synchronization connection is lost, then no state synchronization occurs. If the configuration synchronization is lost, heartbeats are lost. Both devices determine that the other is down, and both become active.
- You can configure the management ports (MGT) on the HA devices to provide a backup path for heartbeat and hello messages. If you configure HA1 or HA1 backup on the management port, you do not have to enable the heartbeat backup option.

## Active/Active HA

Active/active high availability allows both devices in an HA pair to pass traffic concurrently and is deployed primarily in asymmetrically routed environments where App-ID and Content-ID support are required. Layer 7 inspection for App-ID and Content-ID is performed on a single device for each session (that device is known as the session owner). PAN-OS uses packet forwarding (through the HA3 link), where required, to send packets to the designated session owner for processing.

***Note:*** *You must enable jumbo frames on the firewall and on all intermediary networking devices when using the HA3 interface. To enable jumbo frames, select* ***Device > Setup < Session*** *and select the option to* ***Enable Jumbo Frame*** *in the Session Settings section.*

Active/active devices can be deployed with Layer 3 or virtual wire interfaces. In Layer 3 deployments, the scanned packets can be forwarded directly by the session owner after processing. In virtual wire deployments, the scanned packets must be returned to the receiving firewall to preserve the forwarding path. If the session owner receives the packet

initially, the HA3 link is not used. Sessions that do not require App-ID and Content-ID are forwarded directly by the receiving device (even if it is not the session owner) to maximize performance.

To provide flexibility, you can configure Layer 3 interfaces in several ways. Note that it often makes sense to configure Layer 3 interfaces with a static interface IP address in addition to a virtual address (floating IP address or ARP load sharing IP address).

- **Static interface IP**—Layer 3 interfaces should be assigned static IP addresses whenever the firewall will be participating in dynamic routing protocols with neighboring devices. One possible active/active deployment option makes use of dynamic routing protocol cost metrics to force a symmetric path through the HA pair. In this case, all traffic will be symmetric and the efficiency of the active/active pair will be maximized.
- **Floating IP**—This mode is employed when Virtual Router Redundancy Protocol (VRRP)-like functionality is required, such as when an IP address must be available regardless of the state of the HA pair members. It is typical to configure two floating IP addresses on a particular interface such that each firewall owns one. Ownership is assigned to the device ID that has the higher priority. If either firewall fails, the floating IP address will be transitioned to the HA peer.
- **ARP load sharing**—This mode is used to distribute the load of host traffic between the two firewalls using Address Resolution Protocol (ARP).

For a more in-depth discussion of these three options refer to the discussions later in this section.

# Packet Flow

Packet flow works as follows in an active/active configuration:

- The session owner is responsible for all packet processing for App-ID and Content-ID. The session owner can be configured to be (1) the first device that receives a packet for the session or (2) the primary device. If the configuration option is set to “primary device,” all sessions are set up on the primary device.

***Note:*** *Logs passing through an active/active HA pair appear on the device that is designated as the session owner.*

- A single device is selected as the session setup device for all new sessions. This is necessary to avoid possible race conditions that can occur in asymmetrically routed environments. The session setup device is determined by one of the following methods:
  - **IP modulo**—Uses a simple modulus operation on the source IP address to determine which device will set up the session. IP modulo distributes session setup responsibilities to a particular HA device according to the parity of the IP address.
  - **Primary Device**—Session setup always occurs on the primary device.
  - **Hash**—Hashing is used to inject more randomness in the setup device selection process.

- When a new session begins, the receiving firewall either sets up the session or forwards it to the HA peer. The action is determined by the session setup configuration, as described previously. During this time, the session owner (the device responsible for maintaining state for App-ID and Content-ID) is determined according to the configuration.
- If packets arrive at the session owner, the packet is scanned for threats (if configured in security policy) and forwarded according to the device's networking configuration. If packets arrive at the HA peer, a session table lookup identifies that the session is owned by the other device and the packet can be forwarded across HA3 to the session owner. If Layer 7 inspection is not required for the session, the receiving device can simply match the session with an existing session table entry and forward the packet towards its final destination.

# Deployment Options

Active/active HA supports concurrent use of virtual wire and Layer 3 interfaces. All active/active deployment options are supported in IPv6 environments, including IPv6 path monitoring.

### Virtual Wire Deployment

Virtual wire deployments support full asymmetric routing as with other active/active deployments. It is important to note that packets forwarded to the session owner for App-ID and Content-ID inspection must be returned to the receiving firewall to preserve the forwarding path.

### Layer 3-Floating IP Deployment

This deployment option allows for the creation of floating IP addresses that can move between the HA devices when a link failure or device failure occurs. The port that owns the floating IP address responds to ARP requests with a virtual MAC address. Floating IP addresses are recommended when VRRP-like functionality is required. Floating IP addresses can be used in VPN and Network Address Translation (NAT) configurations, allowing for persistent connections when a failure occurs on the device offering those services.

***Note:*** *Floating IP address will utilize a different virtual MAC address when the IP moves between HA devices when a failure occurs.*

### Layer 3-ARP Load-Sharing

ARP load-sharing allows the HA pair to share an IP address and provide gateway services. In this scenario, all hosts are configured with a single gateway IP address. ARP requests for the gateway IP address are responded to by a single device in the pair, according to the source of the ARP request. The device selection algorithm can be tuned to achieve a more even distribution of host traffic between the two firewalls. ARP load-sharing should be used when the firewall and hosts exist on the same broadcast domain. If Layer 3 separation exists, the benefits of ARP load-sharing will be lost.

***Note:*** *You cannot ping or perform any management services on an ARP load-sharing IP address in active/active mode.*

### Layer 3-Route Based Redundancy (Static Interface IPs)

Route based redundancy forces traffic to be symmetric by using routing metrics such as Open Shortest Path First (OSPF) costs on the firewalls and on neighboring devices. Load sharing can be handled by adjusting costs to route traffic through both firewalls. In this case, the IP address assigned to the device interface is pinned down and does not fail over to the HA peer during a failover.

## NAT Considerations

In active/active mode, it is necessary to define an active/active device binding in all NAT rules. Active/active device binding becomes available in the web interface when the HA mode has been changed to active/active. When a new session is created, device binding determines which NAT rules are matched by the firewall (the device binding must include the session owner device to produce a match). Although NAT policy match is performed by the session setup device, NAT rules are evaluated from the perspective of the session owner. The session is translated according to NAT rules that are bound to the session owner device. For device-specific rules, a firewall skips all NAT rules that are not bound to the session owner when the NAT policy match is performed.

For example, suppose device 1 is the session owner and is also responsible for setting up the session. When device 1 attempts to match the session to a NAT rule, it will skip all rules with a device binding of device 0.

NAT device binding options include the following:

- **Device 0 and Device 1**—Translation is performed according to device-specific bindings only if the session owner and the device ID in the NAT rule match. Device-specific NAT rules are commonly used when the two firewalls use unique public IP addresses for translation.
- **Both**—This option allows either device to match new sessions to the NAT rule and is commonly used for destination NAT.
- **Primary**—This option allows only the active-primary device to match new sessions to the NAT rule. This setting is used mainly for inbound static NAT, where only one firewall should respond to ARP requests. Unlike device 0/1 bindings, a primary device binding can move between devices when the primary role is transferred.

The following scenarios apply to active/active NAT deployments.

### Source Translation to Dynamic IP or IP/Port Pool

When source translating to a dynamic IP or dynamic IP/port pool, it is necessary to tie NAT rules to a specific device (Device ID 0 or 1). The IP pools to which the HA devices are translating must not overlap. When a session is established, either device can translate return packets.

In the following example, the sessions owned by device 0 are translated to 1.1.1.1 and the sessions owned by device 1 are translated to 1.1.1.2. In the event of a device failure, sessions from device 0 will continue to be translated to 1.1.1.1 until they have ceased. In this example, it is valid to use floating IP addresses on each of the firewalls if that functionality is required.

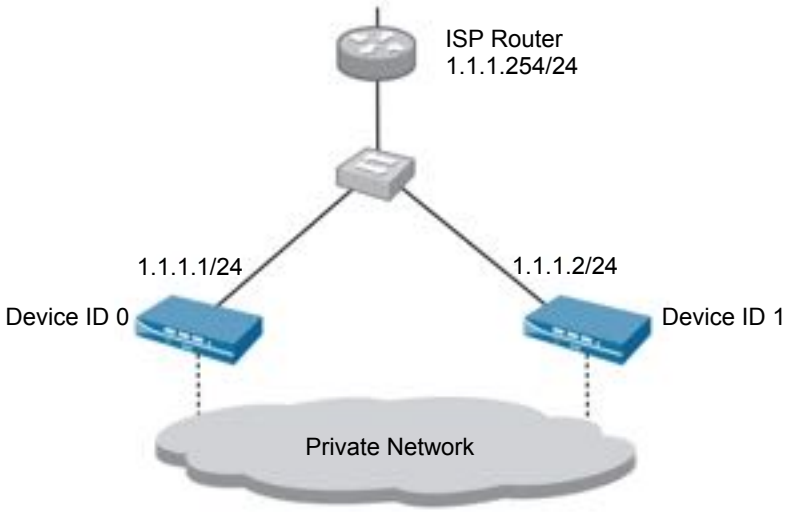


**Figure 3. Source Translation Dynamic IP Configuration**

**Table 40. Source Translation Dynamic IP Rules**

| | Original Packet | | Translated Packet | |
|---|---|---|---|---|
| **Name** | **Source Zone** | **Destination Zone** | **Source Translation** | **Active/Active HA Binding** |
| Src NAT Device 0 | L3Trust | L3Untrust | dynamic-ip-and-port 1.1.1.1 | 0 |
| Src NAT Device 1 | L3Trust | L3Untrust | dynamic-ip-and-port 1.1.1.2 | 1 |

### Dynamic Source Translation to Public IP Addresses for Different Internet Service Providers (ISPs)

In this scenario, NAT rules are tied to specific devices (Device ID 0 or 1). All sessions owned by device 0 are translated to 1.1.1.1 and all sessions owned by device 1 are translated to 2.2.2.1. If device 0 fails, device 1 will attempt to translate existing sessions according to the original IP address of 1.1.1.1. If the second ISP cannot route to these addresses, the sessions will fail. In this example, the ISP-specific interface IP addresses are pinned down to a particular device. A floating IP address should not be used in this configuration.

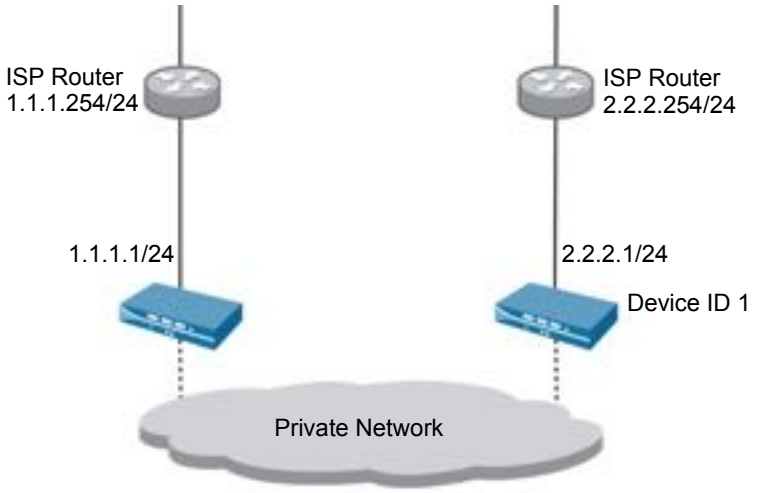


**Figure 4. Dynamic Source Translation to Public IP Address Configuration**

**Table 41. Dynamic Source Translation to Public IP Address Rules**

| | Original Packet | | Translated Packet | |
|---|---|---|---|---|
| **Name** | **Source Zone** | **Destination Zone** | **Source Translation** | **Active/Active HA Binding** |
| Src NAT Device 0 | L3Trust | L3Untrust | dynamic-ip-and-port 1.1.1.1 | 0 |
| Src NAT Device 1 | L3Trust | L3Untrust | dynamic-ip-and-port 2.2.2.1 | 1 |

### Destination Translation to Provider-Independent IP Address

In this scenario, NAT rules are tied to both devices. Translation is the same regardless of which device receives the first incoming packet. A packet destined to 3.3.3.30 will be translated to 10.0.0.200 regardless of which device receives the packet.

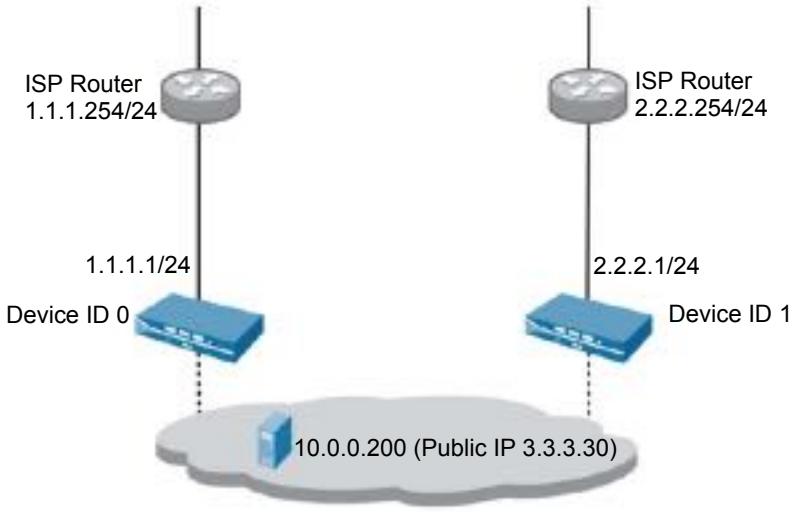


**Figure 5. Destination Translation to Provider-Independent IP Address**

**Table 42. Dynamic Source Translation to Public IP Address Rules**

| | | Original Packet | | Translated Packet | |
|---|---|---|---|---|---|
| **Name** | **Source Zone** | **Destination Zone** | **Destination Address** | **Destination Translation** | **Active/Active HA Binding** |
| DNAT Prov Indep | L3Untrust | L3Untrust | 3.3.3.30 | address: 10.0.0.200 | both |

# Setting Up HA

To set up HA, follow these steps:

1. Use two firewalls with the same model number.
2. Mount both firewalls in a rack near each other, and power them up as described in the *Hardware Reference Guide*. If you are enabling HA on an existing production device, backup the configuration and you may also want to perform a factory reset on the new device that will be introduced. This will ensure that the configuration is clean on the new device. After HA is configured, you would then push the configuration from the existing primary device to the new secondary device. Make sure that you do not push the factory config from the new device to the existing device. For information on performing a factory reset, refer to the *PAN-OS Command Line Interface Reference Guide*.
3. Connect each firewall to your network and the Internet using the same physical ports.

4. Using two crossover RJ-45 Ethernet cables, connect the HA1 and HA2 ports on each firewall to the same ports on the other firewall, or connect the ports on both firewalls to a switch. HA1 is for the control link, and HA2 is for the data link. For active/active configurations, make an additional physical connection, HA3, between the two firewalls. Link aggregation groups are recommended for link redundancy on HA3 when the firewall supports aggregate Ethernet.

***Note:*** *For devices that do not have dedicated HA interfaces, you must use the traffic interfaces for HA. For example, connect the ethernet 1/15 interfaces to each other and the ethernet1/16 interfaces to each other.*

5. Open the **Network** tab and verify that the HA links are up. Configure each to be of the type HA.

**Figure 6. Verifying HA Interfaces**

6. Configure HA settings on both firewalls. Refer to “Enabling HA on the Firewall” on page 101.

### Item to note when setting up HA

Crossover cables are recommended when HA links are directly connected.

# Enabling HA on the Firewall

▶ *Device > High Availability*

After setting up HA as described in "Setting Up HA" on page 99, you can enable HA on both the active and passive firewall. For each section on the **High Availability** page, click **Edit** in the header, and specify the corresponding information described in the following table.

**Table 43. HA Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Setup | Specify the following settings:<br>• **Enable HA**—Activate HA functionality.<br>• **Group ID**—Enter a number to identify the active/passive pair (1 to 63). Allows multiple pairs of active/passive firewalls to reside on the same network. The ID must be unique when more than one high availability pair resides on an Layer 2 network.<br>• **Description**—Enter a description of the active/passive pair (optional).<br>• **Mode**—Choose active-active or active-passive.<br>• **Peer HA IP Address**—Enter the IP address of the HA1 interface that is specified in the Control Link section of the other firewall.<br>• **Backup Peer HA IP Address**—Enter the IP address for the peer's backup control link.<br>• **Enable Config Sync**—Synchronize the peer system.<br>• **Link Speed**— Select the speed for the data link between the active and passive firewalls (Firewalls with dedicated HA ports).<br>• **Link Duplex**— Select a duplex option for the data link between the active and passive firewalls (Firewalls with dedicated HA ports). |

**Table 43. HA Settings (Continued)**

| Field | Description |
|---|---|
| Election Settings | Specify the following settings:<br>• **Device Priority**—Enter a priority value to identify the active firewall. The firewall with the lower value (higher priority) becomes the active firewall (range 0-255).<br>• **Heartbeat Backup**—Uses the management ports on the HA devices to provide a backup path for heartbeat and hello messages. The management port IP address will be shared with the HA peer through the HA1 control link. No additional configuration is required.<br>• **Preemptive**—Enable the higher priority firewall to resume active operation after recovering from a failure. If this setting is off, then the lower priority firewall remains active even after the higher priority firewall recovers from a failure.<br>• **Preemption Hold Time**—Enter the time a passive or active-secondary device will wait before taking over as the active or active-primary device (range 1-60 min, default 1 min).<br>• **Promotion Hold Time**—Enter the time that the passive device (in active/passive mode) or the active-secondary device (in active/active mode) will wait before taking over as the active or active-primary device after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made.<br>• **Hello Interval**—Enter the number of milliseconds between the hello packets sent to verify that the HA program on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms.<br>• **Heartbeat Interval**—Specify how frequently the HA peers exchange heartbeat messages in the form of an ICMP ping (range 1000-60000 ms, default 1000 ms).<br>• **Maximum No. of Flaps**—A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. You can specify the maximum number of flaps that are permitted before the firewall is determined to be suspended and the passive firewall takes over (range 0-16, default 3). The value 0 means there is no maximum (an infinite number of flaps is required before the passive firewall takes over).<br>• **Monitor Fail Hold Up Time (ms)**—Specify the interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is recommended to avoid an HA failover due to the occasional flapping of neighboring devices (range 0-60000 ms, default 0 ms).<br>• **Additional Master Hold Up Time (min)**—This time interval is applied to the same event as Monitor Fail Hold Up Time (range 0-60000 ms, default 500 ms). The additional time interval is applied only to the active device in active/passive mode and to the active-primary device in active/active mode. This timer is recommended to avoid a failover when both devices experience the same link/path monitor failure simultaneously. |

**Table 43. HA Settings (Continued)**

| Field | Description |
|---|---|
| Control Link (HA1)/ Control Link (HA1 Backup) | The recommended configuration for the HA control link connection is to use the dedicated HA1 link between the two devices and use the management port as the Control Link (HA Backup) interface. In this case, you do not need to enable the Heartbeat Backup option in the Elections Settings page. If you are using a physical HA1 port for the Control Link HA link and a data port for Control Link (HA Backup), it is recommended that enable the Heartbeat Backup option.<br>For devices that do not have a dedicated HA port, such as the PA-200, you should configure the management port for the Control Link HA connection and a data port interface configured with type HA for the Control Link HA1 Backup connection. Since the management port is being used in this case, there is no need to enable the Heartbeat Backup option in the Elections Settings page because the heartbeat backups will already occur through the management interface connection.<br>***Note:*** *When using a data port for the HA control link, you should be aware that since the control messages have to communicate from the dataplane to the management plane, if a failure occurs in the dataplane, HA control link information cannot communicate between devices and a failover will occur. It is best to use the dedicated HA ports, or on devices that do not have a dedicated HA port, use the management port.*<br>Specify the following settings for the primary and backup HA control links:<br>• **Port**—Select the HA port for the primary and backup HA1 interfaces. The backup setting is optional.<br>***Note:*** *The management port can also be used as the control link.*<br>• **IPv4/IPv6 Address**—Enter the IPv4 or IPv6 address of the HA1 interface for the primary and backup HA1 interfaces. The backup setting is optional.<br>• **Netmask**—Enter the network mask for the IP address (such as "255.255.255.0") for the primary and backup HA1 interfaces. The backup setting is optional.<br>• **Gateway**—Enter the IP address of the default gateway for the primary and backup HA1 interfaces. The backup setting is optional.<br>• **Link Speed (Models with dedicated HA ports only)**—Select the speed for the control link between the firewalls for the dedicated HA1 port.<br>• **Link Duplex (Models with dedicated HA ports only)**—Select a duplex option for the control link between the firewalls for the dedicated HA1 port.<br>• **Encryption Enabled**—Enable encryption after exporting the HA key from the HA peer and importing it onto this device. The HA key on this device must also be exported from this device and imported on the HA peer. Configure this setting for the primary HA1 interface.<br>The key import/export is done on the Certificates page. Refer to "Importing, Exporting and Generating Security Certificates" on page 60.<br>• **Monitor Hold Time (ms)**—Enter the length of time (milliseconds) that the firewall will wait before declaring a peer failure due to a control link failure (1000-60000 ms, default 3000 ms). This option monitors the physical link status of the HA1 port(s). |

**Table 43. HA Settings (Continued)**

| Field | Description |
|---|---|
| Data Link (HA2) | Specify the following settings for the primary and backup data link:<br>• **Port**—Select the HA port. Configure this setting for the primary and backup HA2 interfaces. The backup setting is optional.<br>• **IP Address**—Specify the IPv4 or IPv6 address of the HA interface for the primary and backup HA2 interfaces. The backup setting is optional.<br>• **Netmask**—Specify the network mask for the HA interface for the primary and backup HA2 interfaces. The backup setting is optional.<br>• **Gateway**—Specify the default gateway for the HA interface for the primary and backup HA2 interfaces. The backup setting is optional. If the HA2 IP addresses of the firewalls in the HA pair are in the same subnet, the Gateway field should be left blank.<br>• **Enable Session Synchronization**—Enable synchronization of the session information with the passive firewall, and choose a transport option.<br>• **Transport**—Choose one of the following transport options:<br>– **Ethernet**—Use when the firewalls are connected back-to-back or through a switch (Ethertype 0x7261).<br>– **IP**—Use when Layer 3 transport is required (IP protocol number 99).<br>– **UDP**—Use to take advantage of the fact that the checksum is calculated on the entire packet rather than just the header, as in the IP option (UDP port 29281).<br>• **Link Speed (Models with dedicated HA ports only)**—Select the speed for the control link between the active and passive firewalls for the dedicated HA2 port.<br>• **Link Duplex (Models with dedicated HA ports only)**—Select a duplex option for the control link between the active and passive firewalls for the dedicated HA2 port.<br>• **HA2 keep-alive**—Select this check box to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set, the defined action will occur (log or split data-path). The option is disabled by default.<br>You can configure the HA2 keep-alive option on both devices, or just one device in the HA pair. If the option is only set on one device, only that device will send the keep-alive messages. The other device will be notified though if a failure occurs and will go into split data path mode if that action is selected.<br>– **Action**—Select the action to take if monitoring messages fail based on the threshold setting.<br>› **log-only**—Select to generate a critical level system log message when an HA2 failure occurs based on the threshold setting. If the HA2 path recovers, an informational log will be generated.<br>In an active/passive configuration, you should use this action since there's no need to split the data since only one device is active at a given time.<br>› **split data-path**—This action is designed for an active/active HA configuration. In an ***active/active*** configuration, if session synchronization is disabled by the admin, or by a monitoring failure, session ownership and session setup will both be set to the local device and new sessions will be processed locally for the session lifetime. |

**Table 43. HA Settings (Continued)**

| Field | Description |
|---|---|
| | › **Threshold (ms)**—The duration in which keep-alive messages have failed before one of the above actions will be triggered (range 5000-60000ms, default 10000ms).<br><br>***Note:*** *When an HA2 backup link is configured, failover to the backup link will occur if there is a physical link failure. With the HA2 keep-alive option enabled, the failover will also occur if the HA keep-alive messages fail based on the defined threshold.* |
| **Link and Path Monitoring Tab** | |
| Path Monitoring | Specify the following:<br>• **Enabled**—Enable path monitoring. Path monitoring enables the firewall to monitor specified destination IP addresses by sending ICMP ping messages to make sure that they are responsive. Use path monitoring for virtual wire, Layer 2, or Layer 3 configurations where monitoring of other network devices is required for failover and link monitoring alone is not sufficient.<br>• **Failure Condition**—Select whether a failover occurs when any or all of the monitored path groups fail to respond. |
| Path Group | Define one or more path groups to monitor specific destination addresses. To add a path group, click **Add** for the interface type (Virtual Wire, VLAN, or Virtual Router) and specify the following:<br>• **Name**—Enter a name to identify the group.<br>• **Enabled**—Enable the path group.<br>• **Failure Condition—**Select whether a failure occurs when any or all of the specified destination addresses fails to respond.<br>• **Source IP**—For virtual wire and VLAN interfaces, enter the source IP address used in the probe packets sent to the next-hop router (Destination IP address). The local router must be able to route the address to the firewall. The source IP address for path groups associated with virtual routers will be automatically configured as the interface IP address that is indicated in the route table as the egress interface for the specified destination IP address.<br>• **Destination IPs**—Enter one or more (comma-separated) destination addresses to be monitored.<br>• **Ping Interval**—Specify the interval between pings that are sent to the destination address (range 200-60,000 milliseconds, default 200 milliseconds).<br>• **Ping Count**—Specify the number of failed pings before declaring a failure (range 3-10 pings, default 10 pings). |

**Table 43. HA Settings (Continued)**

| Field | Description |
|---|---|
| Link Monitoring | Specify the following:<br>• **Enabled**—Enable link monitoring. Link monitoring allows failover to be triggered when a physical link or group of physical links fails.<br>• **Failure Condition**—Select whether a failover occurs when any or all of the monitored link groups fail. |
| Link Groups | Define one or more link groups to monitor specific Ethernet links. To add a link group, specify the following and click **Add**:<br>• **Name**—Enter a link group name.<br>• **Enabled**—Enable the link group.<br>• **Failure Condition—**Select whether a failure occurs when any or all of the selected links fail.<br>• **Interfaces**—Select one or more Ethernet interfaces to be monitored. |
| **Active/Passive Tab** | |
| Passive Link State | Choose from the following options:<br>• **auto**—Causes the link status to reflect physical connectivity, but discards all packets received. This option allows the link state of the interface to stay up until a failover occurs, decreasing the amount of time it takes for the passive device to take over.<br>This option is supported in Layer 2, Layer 3, and Virtual Wire mode. The auto option is desirable, if it is feasible for your network.<br>***Note:*** *When set to auto, the link state on the passive peer in a VM-Series firewall configured as an HA pair displays in a down state on* ***Network > Interfaces > Ethernet****. The link state for the interface icon is red.*<br>• **shutdown**—Forces the interface link to the down state. This is the default option, which ensures that loops are not created in the network. |
| Monitor Fail Hold Down Time | Specify the length of time (minutes) that a firewall will spend in the non-functional state before becoming passive. This timer is used only when the failure reason is a link or path monitor failure (range 1 to 60, default 1). |
| **Active/Active Config Tab** | |
| Packet Forwarding | Select the Enable check box to enable forwarding of packets over the HA3 link. This is required for asymmetrically routed sessions that require Layer 7 inspection for App-ID and Content-ID. |
| HA3 Interface | Choose the interface to forward packets between HA peers when configured in active/active mode.<br>***Note:*** *You must enable jumbo frames on the firewall and on all intermediary networking devices when using the HA3 interface. To enable jumbo frames, select* ***Device > Setup < Session*** *and select the option to* ***Enable Jumbo Frame*** *in the Session Settings section.* |
| VR Sync | Force synchronization of all virtual routers configured on the HA devices.<br>Virtual Router synchronization can be used when the virtual router is not employing dynamic routing protocols. Both devices must be connected to the same next-hop router through a switched network and must use only static routing. |

**Table 43. HA Settings (Continued)**

| Field | Description |
|---|---|
| QoS Sync | Synchronize the QoS profile selection on all physical interfaces. Use this option when both devices have similar link speeds and require the same QoS profiles on all physical interfaces. This setting affects the synchronization of QoS settings on the **Network** tab. QoS policy is synchronized regardless of this setting. |
| Tentative Hold Time (sec) | When a firewall in an HA active/active state fails it will go into a tentative state. This timer defines how long it will stay in this state. During the tentative period the firewall will attempt to build routing adjacencies and populate its route table before it will process any packets. Without this timer, the recovering firewall would enter the active-secondary state immediately and would blackhole packets because it would not have the necessary routes (default 60 seconds). |
| Session Owner Selection | Specify one of the following options for selecting the session owner:<br>• **Primary Device**—Select this option to have the active/primary firewall handle Layer 7 inspection for all sessions. This setting is recommended primarily for troubleshooting operations.<br>• **First packet**—Select this option to make the firewall that receives the first packet of the session responsible for Layer 7 inspection in support of App-ID and Content-ID. This is the recommended configuration to minimize utilization of the HA3 packet forwarding link. |
| Session Setup | Choose the method for initial session setup.<br>• **IP Modulo**—Selects a firewall based on the parity of the source IP address.<br>• **Primary Device**—Ensures that all sessions are set up on the primary firewall.<br>• **IP Hash**—Determines the setup firewall using a hash of the source IP address or source and destination IP address, and hash seed value if more randomization is desired. |

**Table 43. HA Settings (Continued)**

| Field | Description |
|---|---|
| Virtual Address | Click **Add**, select the IPv4 or IPv6 tab and then click **Add** again to enter options for an HA virtual Address that will be used by the HA active/active cluster. You can select the type of virtual address to be either Floating or ARP Load Sharing. You can also mix the type of virtual address types in the cluster, for example, you could use ARP load sharing on the LAN interface and a Floating IP on the WAN interface. For more information deployment types, refer to "Deployment Options" on page 95.<br>• **Floating**—Enter an IP address that will move between HA devices in the event of a link or device failure. You should configure two floating IP addresses on the interface, so that each firewall will own one and then set the priority. If either firewall fails, the floating IP address will be transitioned to the HA peer.<br>– Device 0 Priority—Set the priority to determine which device will own the floating IP address. A device with the lowest value will have the highest priority.<br>– Device 1 Priority—Set the priority to determine which device will own the floating IP address. A device with the lowest value will have the highest priority.<br>– Failover address if link state is down—Use the failover address when the link state is down on the interface.<br>• **ARP Load Sharing**—Enter an IP address that will be shared by the HA pair and will provide gateway services for hosts. This option should only be used when the firewall and hosts exist on the same broadcast domain. Select the **Device Selection Algorithm**:<br>– **IP Modulo**—If this option is selected, the firewall that will respond to ARP requests will be selected based on the parity of the ARP requesters IP address.<br>– **IP Hash**—If this option is selected, the firewall that will respond to ARP requests will be selected based on a hash of the ARP requesters IP address. |
| **Operational Commands** | |
| Suspend local device | Put the HA device in suspend mode, which temporarily disables HA functionality on the firewall. If you suspend the current active firewall, the secondary firewall will take over. |

### Important items to consider when configuring HA

- The firewall that is assigned the lower device priority value is the higher priority device and becomes the active firewall in an HA pair when preemption is enabled on both firewalls in the pair.
- The Preemption option must be enabled on both devices for the higher priority firewall to resume active operation upon recovery following a failure.
- The subnet that is used for the local and peer IP should not be used anywhere else on the virtual router.
- The OS and Content versions should be the same on each device. A mismatch can prevent the devices in the cluster from synchronizing.
- The LEDs are green on the HA ports for the active firewall and amber on the passive firewall.

- To test failover, pull a cable on the active device, or put the active device into a suspend state by issuing the CLI command `request high-availability state suspend`. You can also suspend the active device by clicking the **Suspend** link at **Device > High Availability > Operational Commands** tab**.**
- To place a suspended device back into a functional state, use the CLI command `request high-availability state functional`.
- To view detailed HA information about the local firewall, use the CLI command **show high-availability all**.
- To compare the configuration of the local and peer firewalls, use the CLI command **show high-availability state** from either device. You can also compare the configurations on the local and peer firewalls using the **Config Audit** tool on the **Device** tab by selecting the desired local configuration in the left selection box and the peer configuration in the right selection box.
- Synchronize the firewalls from the web interface by pressing the **Push Configuration** button located in the HA widget on the **Dashboard** tab. Note that the configuration on the device from which you push the configuration overwrites the configuration on the peer device. To synchronize the firewalls from the CLI on the active device, use the command `request high-availability sync-to-remote running-config`.
- To follow the status of the load, use the CLI command **show jobs processed**.

***Note:*** *In a High Availability (HA) active/passive configuration with devices that use 10 Gigabit SFP+ ports, when a failover occurs and the active device changes to a passive state, the 10 Gigabit Ethernet port is taken down and then brought back up to refresh the port, but does not enable transmit until the device becomes active again. If you have monitoring software on the neighboring device, it will see the port as flapping because it is going down and then up again. This is different behavior than the action with other ports, such as the 1 Gigabit Ethernet port, which is disabled and still allows transmit, so flapping is not detected by the neighboring device.*

### HA Lite

The PA-200 and VM-Series firewalls supports a “lite” version of active/passive HA that does not include any session synchronization. HA lite does provide configuration synchronization and synchronization of some runtime items. It also supports failover of IPSec tunnels (sessions must be re-established), DHCP server lease information, DHCP client lease information, PPPoE lease information, and the firewall's forwarding table when configured in Layer 3 mode.

# Virtual Systems

A virtual system specifies a collection of physical and logical firewall interfaces (including VLANs, and virtual wires) and security zones. (For more information on security zones, refer to "Defining Security Zones" on page 151.) Virtual systems allow you to segment the administration of all policies (security, NAT, QoS, etc.) as well as all reporting and visibility functions provided by the firewall.

Virtual systems generally operate on the security functionality of the firewall. Networking functions including static and dynamic routing are not controlled by virtual systems. If routing segmentation is desired for each virtual system, you must create an additional virtual router.

***Note:*** *The PA-4000 and PA-5000 Series firewalls support multiple virtual systems. The PA-2000 and PA-3000 Series firewalls can support multiple virtual systems if the appropriate license is installed. The PA-500 and PA-200 firewalls do not support virtual systems.*

For example, if you want to customize the security features for the traffic that is associated with your Finance department, you can define a Finance virtual system and then define security policies to apply only to that department.

Figure 7 illustrates the relationship between policies and virtual systems in the firewall. Policies are associated with individual virtual systems, by contrast with device and network level functions, which apply to the overall firewall.

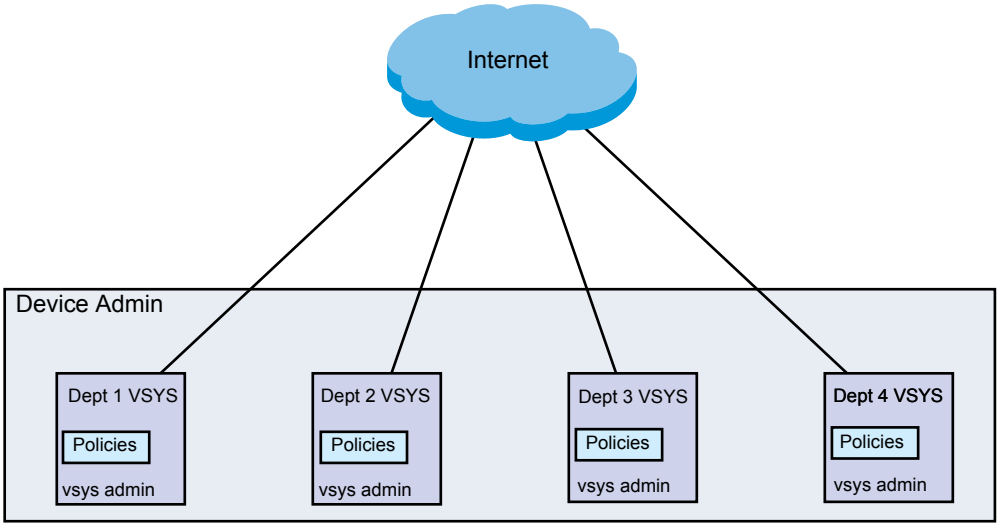


**Figure 7. Virtual Systems and Policies**

To optimize policy administration, you can create virtual system administrator accounts that allow access to individual virtual systems, while maintaining separate administrator accounts for overall device and network functions. For example, a virtual system administrator in the Finance department can be assigned to manage the security policies only for that department.

Initially all interfaces, zones, and policies belong to the default virtual system (vsys1).

When you enable multiple virtual systems, note the following:

- All items needed for policies are created and administered by a virtual systems administrator.
- Zones are objects within virtual systems. Before defining a policy or policy object, select the virtual system from the **Virtual System** drop-down list on the **Policies** or **Objects** tab.
- Interfaces, VLANs, virtual wires, and virtual routers can be assigned to virtual systems. Refer to "Defining Virtual Systems" on page 113.
- Remote logging destinations (SNMP, syslog, and email), as well as applications, services, and profiles, can be shared by all virtual systems or be limited to a selected virtual system.

## Communications Among Virtual Systems

The virtual systems in the firewall are treated as separate entities. To support internal traffic flows between virtual systems, you must indicate which virtual systems are able to communicate with each other. You do so when configuring a virtual system by specifying the other virtual systems that are visible to it. When the virtual systems are made visible to each other, create "external"-type zones and specify which virtual systems will map to each external zone. In the following example, Dept 1 VSYS must have an external zone that refers to Dept 2 VSYS for use in all policies affecting traffic passing between the virtual systems. Traffic log entries are recorded in both virtual systems for inter-VSYS traffic.

Each virtual system must have policies for sending and receiving traffic. For example, allowing Dept 1 VSYS to communicate with Dept 2 VSYS requires a policy in Dept 1 VSYS to allow traffic to go to Dept 2 VSYS and a policy in Dept 2 VSYS to accept incoming traffic from Dept 1 VSYS.

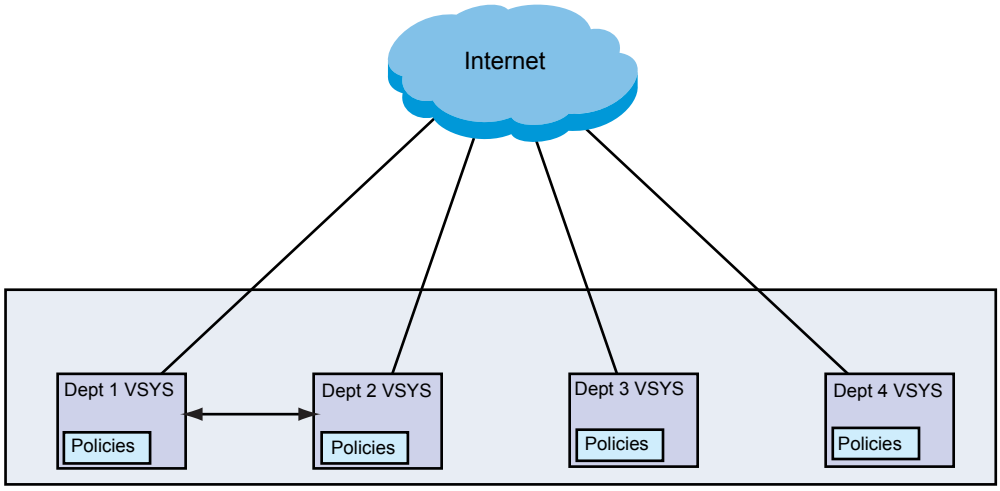


**Figure 8. Communications Among Virtual Systems**

# Shared Gateways

In a standard virtual system interface configuration, each virtual system uses a dedicated interface to the outside world. Each virtual system is autonomous, and there are no direct communication paths among the virtual systems that are internal to the firewall, unless such communications are explicitly configured (refer to "Communications Among Virtual Systems" on page 111). Because each virtual system has its own IP address, multiple addresses are required for external communications.

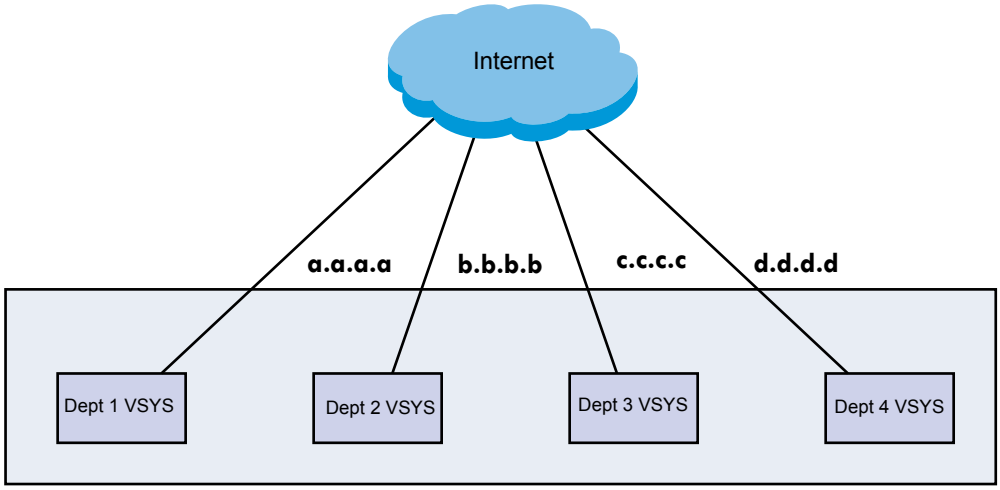


**Figure 9. Virtual Systems Without a Shared Gateway**

Shared gateways allow virtual systems to share a common interface for external communications. This is especially helpful in deployments where the ISP provides only a single IP address. All of the virtual systems communicate with the outside world through the physical interface using a single IP address (see Figure 10). A single virtual router is used to route the traffic for all of the virtual systems through the shared gateway.

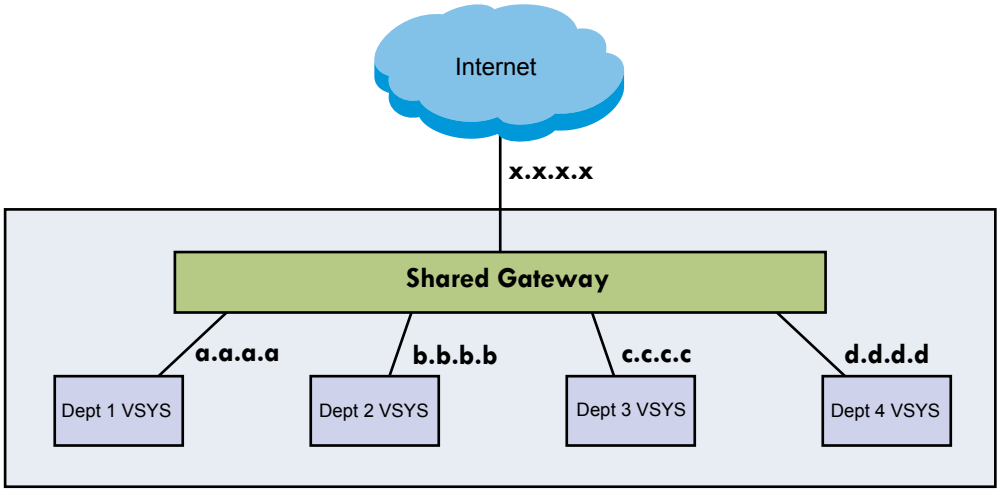


**Figure 10. Virtual Systems with a Shared Gateway**

All policy rules are managed at the virtual system level. You can create NAT and policy-based forwarding rules through the shared gateway, if needed, by selecting the shared gateway from the Virtual System drop-down list on the policy screen.

## Defining Virtual Systems

▶ *Device > Virtual Systems*

To define virtual systems, you must first enable the definition of multiple virtual systems. To do so, open the **Device > Setup** page, click the **Edit** link under **General Settings** in the **Management** tab, and select the **Multi Virtual System Capability** check box. This adds a **Virtual Systems** link to the side menu.

You can now open the **Virtual Systems** page, click **Add**, and specify the following information.

**Table 44. Virtual System Settings**

| Field | Description |
|---|---|
| ID | Enter an integer identifier for the virtual system. Refer to the data sheet for your firewall model for information on the number of supported virtual systems. |
| Name | Enter a name (up to 31 characters) to identify the virtual system. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required. |
| General Tab | Select a DNS proxy profile from the drop-down list if you want to apply DNS proxy rules to this interface. Refer to "DNS Proxy" on page 173.<br>To include objects of a particular type, select the check box for that type (interface, VLAN, virtual wire, virtual router, or visible virtual system). Click **Add** and choose from the drop-down list. You can add one or more objects of any type. To remove an object, select it and click **Delete**. |

**Table 44. Virtual System Settings (Continued)**

| Field | Description |
|---|---|
| Resource Tab | Enter the following settings:<br>• **Sessions Limit**—Maximum number of sessions allowed for this virtual system.<br>• **Security Rules**—Maximum number of security rules allowed for this virtual system.<br>• **NAT Rules**—Maximum number of NAT rules allowed for this virtual system.<br>• **Decryption Rules**—Maximum number decryption rules allowed for this virtual system.<br>• **QoS Rules**—Maximum number of QoS rules allowed for this virtual system.<br>• **Application Override Rules**—Maximum number of application override rules allowed for this virtual system.<br>• **PBF Rules**—Maximum number of policy based forwarding (PBF) rules allowed for this virtual system.<br>• **CP Rules**—Maximum number of captive portal (CP) rules allowed for this virtual system.<br>• **DoS Rules** —Maximum number of denial of service (DoS) rules allowed for this virtual system.<br>• **Site to Site VPN Tunnels**—Maximum number of site-to-site VPN tunnels allowed for this virtual system.<br>• **Concurrent GlobalProtect Tunnel Mode Users**—Maximum number of concurrent remote GlobalProtect users allowed for this virtual system. |

After defining the virtual systems, you can perform any of the following additional tasks:

- To change a virtual system, click the virtual system name or the name of the interface, VLAN, virtual wire, virtual router, or visible virtual systems you want to change, make the appropriate changes, and click **OK**.
- To define security zones for the new virtual system, choose **Network > Zones** and define security zones for each new virtual system (refer to "Defining Security Zones" on page 151). When you define a new zone, you can now select a virtual system.
- Click **Network > Interfaces** and verify that each interface has a virtual system and security zone.

## Configuring Shared Gateways

▶ *Device > Shared Gateways*

Shared gateways use Layer 3 interfaces, and at least one Layer 3 interface must be configured as a shared gateway. Refer to "Configuring Layer 3 Interfaces" on page 130.

**Table 45. Shared Gateway Settings**

| Field | Description |
|---|---|
| ID | Identifier for the gateway (not used by firewall). |
| Name | Enter a name for the shared gateway (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required. |
| Interfaces | Select check boxes for the interfaces that the shared gateway will use. |

# Defining Custom Response Pages

▶ *Device > Response Pages*

Custom response pages are the web pages that are displayed when a user tries to access a URL. You can provide a custom HTML message that is downloaded and displayed instead of the requested web page or file.

Each virtual system can have its own custom response pages.

The following table describes the types of custom response pages that support customer messages.

***Note:*** *Refer to* ***Appendix A, "Custom Pages"*** *for examples of the default response pages.*

**Table 46. Custom Response Page Types**

| Page Type | Description |
|---|---|
| Antivirus Block | Access blocked due to a virus infection. |
| Application Block | Access blocked because the application is blocked by a security policy. |
| Captive Portal Comfort | Page for users to verify their user name and password for machines that are not part of the domain. |
| File Blocking Block | Access blocked because access to the file is blocked. |
| File Blocking Continue | Page for users to confirm that downloading should continue. This option is available only if continue functionality is enabled in the security profile. Refer to "File Blocking Profiles" on page 220. |
| GlobalProtect Portal Help | Custom help page for GlobalProtect users (accessible from the portal). |
| GlobalProtect Portal Login | Page for users who attempt to access the GlobalProtect portal. For information on GlobalProtect, refer to "Overview" on page 335. |

**Table 46. Custom Response Page Types (Continued)**

| Page Type | Description |
|---|---|
| GlobalProtect Welcome Page | Welcome page for users who attempt to log in to the GlobalProtect portal. For information on GlobalProtect, refer to "Overview" on page 335. |
| SSL Certificate Errors Notify Page | Notification that an SSL certificate has been revoked. |
| SSL Decryption Opt-out Page | User warning page indicating that this session will be inspected. |
| URL Filtering Continue and Override Page | Page with initial block policy that allows users to bypass the block. For example, a user who thinks the page was blocked inappropriately can click the **Continue** button to proceed to the page.<br>With the override page, a password is required for the user to override the policy that blocks this URL. See the "URL Admin Override" section of Table 1 for instructions on setting the override password. |
| URL Filtering and Category Match Block Page | Access blocked by a URL filtering profile or because the URL category is blocked by a security policy. |

You can perform any of the following functions under **Response Pages**.

- To import a custom HTML response page, click the link of the page type you would like to change and then click import/export. Browse to locate the page. A message is displayed to indicate whether the import succeeded. For the import to be successful, the file must be in HTML format.
- To export a custom HTML response page, click the **Export** link for the type of page. Select whether to open the file or save it to disk, and select the check box if you want to always use the same option.
- To enable or disable the **Application Block** page or **SSL Decryption Opt-out** pages, click the **Enable** link for the type of page. Select or deselect the **Enable** check box.
- To use the default response page instead of a previously uploaded custom page, delete the custom block page and commit. This will set the default block page as the new active page.

# Viewing Support Information

▶ *Device > Support*

The support page allows you to access support related options. You can view the Palo Alto Networks contact information, view your support expiration date, and view product and security alerts from Palo Alto Networks based on the serial number of your firewall.

Perform any of the following functions on this page:

- **Support**—Use this section to view Palo Alto Networks support contact information, view support status for the device or activate your contract using an authorization code.
- **Production Alerts/Application and Threat Alerts**—These alerts will be retrieved from the Palo Alto Networks update servers when this page is accessed/refreshed. To view the details of production alerts, or application and threat alerts, click the alert name. Production alerts will be posted if there is a large scale recall or urgent issue related to a given release. The application and threat alerts will be posted if significant threats are discovered.
- **Links**—This section provides a link to the Support home page, the site to manage your cases, and a link to register the device using your support login.
- **Tech Support File**—Use the **Generate Tech Support File** link to generate a system file that the Support group can use to help troubleshoot issues that you may be experiencing with the device. After you generate the file, click **Download Tech Support File** to retrieve it and then send it to the Palo Alto Networks Support department.
- **Stats Dump File**—Use the **Generate Stats Dump File** link to generate a set of XML reports that summarizes network traffic over the last 7 days. Once the report is generated, click the **Download Stats Dump File** link to retrieve the report. The files are then used by a Palo Alto Networks or Authorized Partner systems engineer to generate an Application Visibility and Risk Report (AVR Report). The AVR highlights what has been found on the network and the associated business or security risks that may be present and is typically used as part of the evaluation process. For more information on the AVR Report, please contact you Palo Alto Networks or Authorized Partner systems engineer.

## Chapter 4
# Network Configuration

This chapter describes how to configure the firewall to support your network architecture:

- "Firewall Deployment" in the next section
- "Firewall Interfaces" on page 127
- "Security Zones" on page 151
- "VLAN Support" on page 152
- "Virtual Routers and Routing Protocols" on page 153
- "DHCP Server and Relay" on page 171
- "DNS Proxy" on page 173
- "Network Profiles" on page 174

***Note:*** *For information about VPN support on the firewall, refer to "Configuring IPSec Tunnels" on page 309 and "Configuring IPSec Tunnels" on page 309. For information about quality of service (QoS) support, refer to "Configuring Quality of Service" on page 355*

# Firewall Deployment

The firewall can replace your existing firewall when installed between an edge router (or other device that faces the Internet) and a switch or router that connects to your internal network. The firewall supports a wide range of deployment options and interface types that can be used simultaneously on different physical interfaces. They are described in the following sections:


- "Virtual Wire Deployments" in the next section
- "Layer 2 Deployments" on page 124
- "Layer 3 Deployments" on page 124
- "Tap Mode Deployments" on page 125
- "Defining Virtual Wires" on page 125
- "Packet Content Modification" on page 126


## Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together (Figure 11), and should be used only when no switching or routing is needed.

A virtual wire deployment allows the following conveniences:

- Simplifies installation and configuration.
- Does not require any configuration changes to surrounding or adjacent network devices.

The "default-vwire" that is shipped as the factory default configuration, binds together Ethernet ports 1 and 2 and allows all untagged traffic. You can, however, use a virtual wire to connect any two ports and configure it to block or allow traffic based on the virtual LAN (VLAN) tags; the VLAN tag "0" indicates untagged traffic. You can also create multiple subinterfaces, add them into different zones and then classify traffic according to a VLAN tag, or a combination of a VLAN tag with IP classifiers (address, range, or subnet) to apply granular policy control for specific VLAN tags or for VLAN tags from a specific source IP address, range, or subnet.

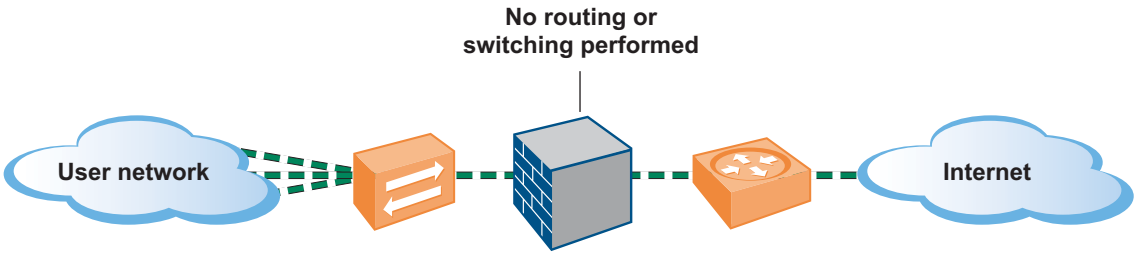


**Figure 11. Virtual Wire Deployment**

### Virtual Wire Subinterfaces

Virtual wire subinterfaces provide flexibility in enforcing distinct policies when you need to manage traffic from multiple customer networks. It allows you to separate and classify traffic into different zones (the zones can belong to separate virtual systems, if required) using the following criteria:

- VLAN tags —The example in Figure 12, shows an Internet Service Provider (ISP) using virtual wire subinterfaces with VLAN tags to separate traffic for two different customers.
- VLAN tags in conjunction with IP classifiers (address, range, or subnet). — The example in Figure 13, shows an Internet Service Provider (ISP) with two separate virtual systems on a firewall that manages traffic from two different customers. On each virtual system, the example illustrates how virtual wire subinterfaces with VLAN tags and IP classifiers are used to classify traffic into separate zones and apply relevant policy for customers from each network.

The general workflow for the configuration is:

1. Configure two Ethernet interfaces as type virtual wire, and assign these interfaces to a virtual wire. To set up virtual wires, see "Configuring Virtual Wire Interfaces" on page 138.
2. Create subinterfaces on the parent Virtual Wire to separate CustomerA and CustomerB traffic. Make sure that the VLAN tags defined on each pair of subinterfaces that are configured as virtual wire(s) are identical. This is essential because a virtual wire does not switch VLAN tags.
   To set up virtual wire subinterfaces, see "Configuring Virtual Wire Subinterfaces" on page 139.
3. Create new subinterface(s) and define IP classifiers. This task is optional and only required if you wish to add additional subinterfaces with IP classifiers for further managing traffic from a customer based on the combination of VLAN tags and a specific source IP address, range or subnet.

   When creating subinterfaces with VLAN tags and IP classifiers you must first create a subinterface with the same VLAN tag but without any IP classifiers and then create subinterfaces with the IP classifiers. Doing so allows traffic that doesn't match the IP classifier to be assigned to the first interface you created.

   You can also use IP classifiers for managing untagged traffic. To do so, you must create a sub-interface with the vlan tag "0", and define sub-interface(s) with IP classifiers for managing untagged traffic using IP classifiers.

***Note:*** *IP classification may only be used on the subinterfaces associated with one side of the virtual wire. The subinterfaces defined on the corresponding side of the virtual wire must use the same VLAN tag, but must not include an IP classifier.*

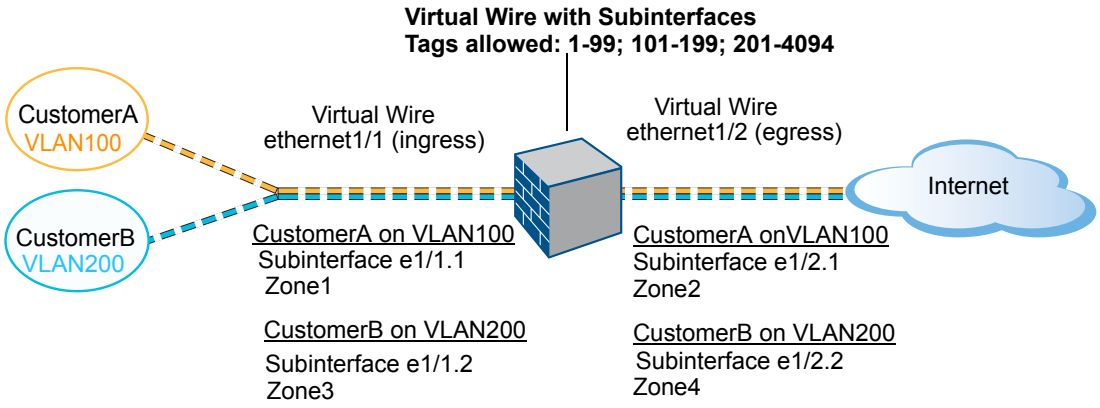


**Figure 12. Virtual Wire Deployment with Subinterfaces (VLAN Tags only)**

The illustration in Figure 12 depicts CustomerA and CustomerB connected to the firewall through one physical interface, ethernet1/1, configured as a Virtual Wire; it is the ingress interface. A second physical interface, ethernet1/2, is also part of the Virtual Wire; it is the egress interface that provides access to the Internet. For CustomerA, you also have subinterfaces ethernet1/1.1 (ingress) and ethernet1/2.1 (egress). For CustomerB, you have the subinterface ethernet1/1.2 (ingress) and ethernet1/2.2 (egress). When configuring the subinterfaces, you must assign the appropriate VLAN tag and zone in order to apply policies for each customer. In this example, the policies for CustomerA are created between Zone1 and Zone2, and policies for CustomerB are created between Zone3 and Zone4.

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this example, a single subinterface matches the VLAN tag on the incoming packet, hence that subinterface is selected. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

***Note:*** *The same VLAN tag must not be defined on the parent virtual wire interface and the subinterface. Verify that the VLAN tags defined on the* ***Tag Allowed*** *list of the parent virtual wire interface (****Network > Virtual Wires****) are not included on a subinterface.*

The illustration in Figure 13 depicts CustomerA and CustomerB connected to one physical firewall that has two virtual systems (vsys), in addition to the default virtual system (vsys1). Each virtual system is an independent virtual firewall that is managed separately for each customer. Each vsys has attached interfaces/subinterfaces and security zones that are managed independently.

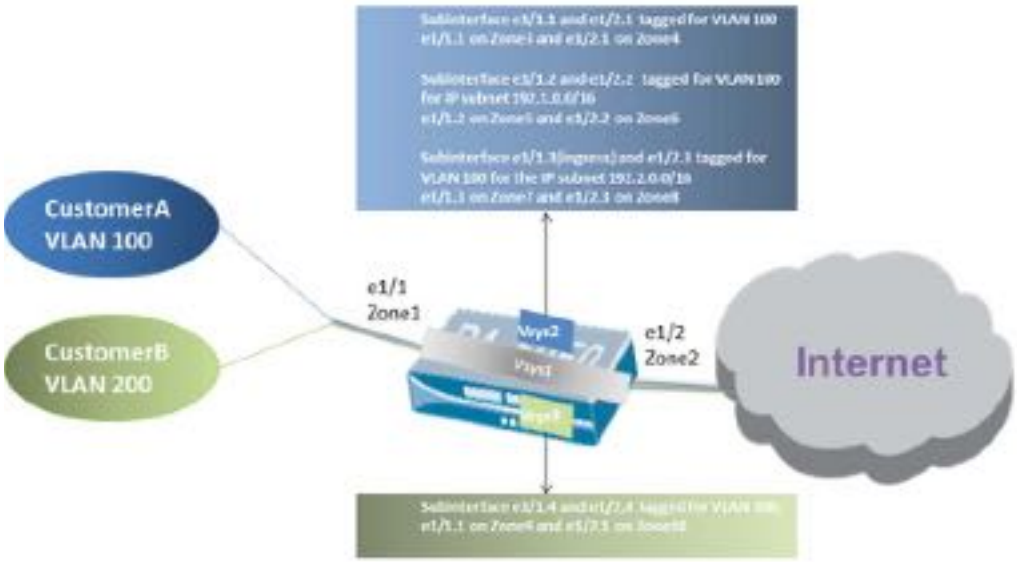


**Figure 13. Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP classifiers)**

Vsys1 is set up to use the physical interfaces ethernet1/1 and ethernet1/2 as a virtual wire; ethernet1/1 is the ingress interface and ethernet1/2 is the egress interface that provides access to the Internet. This virtual wire is configured to accept all tagged and untagged traffic with the exception of VLAN tags 100 and 200 that are assigned to the subinterfaces.

CustomerA is managed on vsys2 and CustomerB is managed on vsys3. On vsys2 and vsys3, the following vwire subinterfaces are created with the appropriate VLAN tags and zones to enforce policy measures:

| Customer | Vsys | Vwire Subinterfaces | Zone | VLAN Tag | IP Classifier |
|---|---|---|---|---|---|
| A | 2 | e1/1.1 (ingress)<br>e1/2.1 (egress) | Zone3<br>Zone4 | 100<br>100 | None |
| | 2 | e1/1.2 (ingress)<br>e1/2.2 (egress) | Zone5<br>Zone6 | 100<br>100 | IP subnet 192.1.0.0/16 |
| | 2 | e1/1.3 (ingress)<br>e1/2.3 (egress) | Zone7<br>Zone8 | 100<br>100 | IP subnet 192.2.0.0/16 |
| B | 3 | e1/1.4 (ingress)<br>e1/2.4 (egress) | Zone9<br>Zone10 | 200<br>200 | None |

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this case, for CustomerA, there are multiple subinterfaces that use the same VLAN tag. Hence, the firewall first narrows the classification to a subinterface based on the source IP address in the packet. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

For return-path traffic, the firewall compares the destination IP address as defined in the IP classifier on the customer-facing subinterface and selects the appropriate virtual wire to route traffic through the accurate subinterface.

***Note:*** *When a session(s) created on a subinterface with VLAN and IP classifiers is offloaded to the hardware for forwarding, the interface counter on the VLAN interface is incremented instead of the subinterface with the IP classifiers. The session counters, however, are accurate for all the interfaces and subinterfaces.*

## Layer 2 Deployments

In a Layer 2 deployment, the firewall provides switching between two or more networks. Each group of interfaces must be assigned to a VLAN object in order for the firewall to switch between them. The firewall will perform VLAN tag switching when layer 2 subinterfaces are attached to a common VLAN object. Choose this option when switching is required (Figure 14).

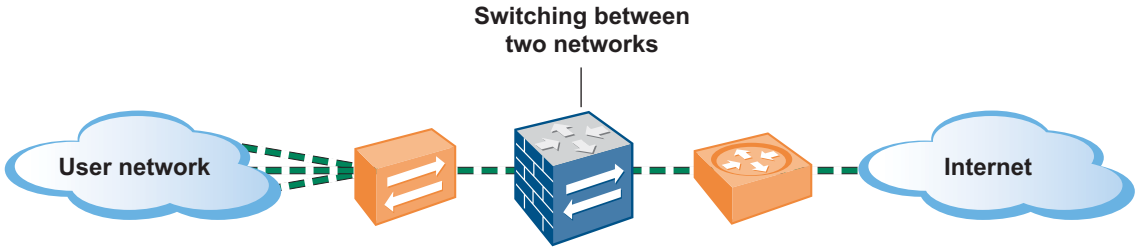


**Figure 14. Layer 2 Deployment**

## Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between multiple ports. An IP address must be assigned to each interface and a virtual router must be defined to route the traffic. Choose this option when routing is required (Figure 15).

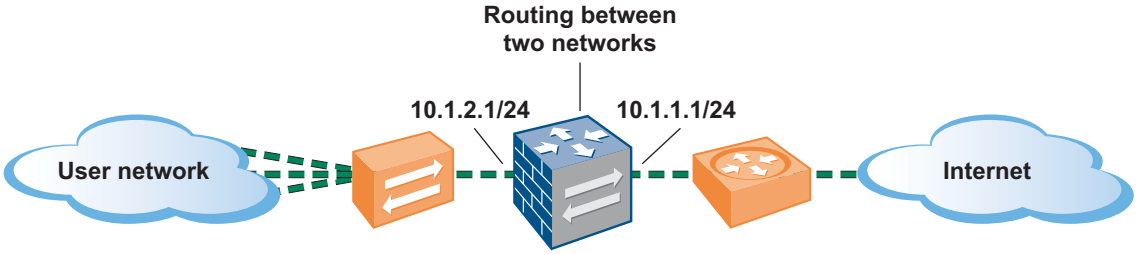


**Figure 15. Layer 3 Deployment**

### Point-to-Point Protocol over Ethernet Support

You can configure the firewall to be a Point-to-Point Protocol over Ethernet (PPPoE) termination point to support connectivity in a Digital Subscriber Line (DSL) environment where there is a DSL modem but no other PPPoE device to terminate the connection.

You can choose the PPPoE option and configure the associated settings when an interface is defined as a Layer 3 interface. For instructions, refer to “Configuring Layer 3 Interfaces” on page 130.

***Note:*** *PPPoE is not supported in HA active/active mode.*

### DHCP Client

You can configure the firewall interface to act as a DHCP client and receive a dynamically assigned IP address. The firewall also provides the capability to propagate settings received by the DHCP client interface into a DHCP server operating on the firewall. This is most commonly used to propagate DNS server settings from an Internet service provider to client machines operating on the network protected by the firewall.

***Note:*** *DHCP client is not supported in HA active/active mode.*

## Tap Mode Deployments

A network tap is a device that provides a way to access data flowing across a computer network. Tap mode deployment allows you to passively monitor traffic flows across a network by way of a switch SPAN or mirror port.

The SPAN or mirror port permits the copying of traffic from other ports on the switch. By dedicating an interface on the firewall as a tap mode interface and connecting it with a switch SPAN port, the switch SPAN port provides the firewall with the mirrored traffic. This provides application visibility within the network without being in the flow of network traffic.

***Note:*** *When deployed in tap mode, the firewall is not able to take action, such as blocking traffic or applying QoS traffic control.*

## Defining Virtual Wires

▶ ***Network > Virtual Wires***

Use this page to define virtual wires after you have specified two virtual wire interfaces on the firewall. For an overview of virtual wire deployments, refer to "Virtual Wire Deployments" on page 120. For instructions on specifying interfaces as virtual wire, refer to "Configuring Virtual Wire Interfaces" on page 138.

**Table 47. Virtual Wire Settings**

| Field | Description |
|---|---|
| Virtual Wire Name | Enter a virtual wire name (up to 31 characters). This name appears in the list of virtual wires when configuring interfaces. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Interfaces | Select two Ethernet interfaces from the displayed list for the virtual wire configuration. Interfaces are listed here only if they have the virtual wire interface type and have not been assigned to another virtual wire. |

**Table 47. Virtual Wire Settings (Continued)**

| Field | Description |
|---|---|
| Tags Allowed | Enter the tag number (0 to 4094) or range of tag numbers (tag1-tag2) for the traffic allowed on the virtual wire. A tag value of zero indicates untagged traffic (the default). Multiple tags or ranges must be separated by commas. Traffic that has an excluded tag value is dropped. Note that tag values are not changed on incoming or outgoing packets.<br>When utilizing virtual wire subinterfaces, the **Tag Allowed** list will cause all traffic with the listed tags to be classified to the parent virtual wire. Virtual wire subinterfaces must utilize tags that do not exist in the parent's **Tag Allowed** list. |
| Multicast Firewalling | Select this option if you want to be able to apply security rules to multicast traffic. If this setting is not enabled, multicast traffic is forwarded across the virtual wire. |
| Link State Pass Through | Select this check box if you want to bring down the other port in a virtual wire when a down link state is detected. If this check box is not selected, link status is not propagated across the virtual wire. |

To change a virtual wire name or the allowed tags, click the virtual wire name on the **Virtual Wires** page, change the settings, and click **OK**. Virtual wires also can be changed from the **Interfaces** page (refer to "Configuring Virtual Wire Interfaces" on page 138).

To delete one or more virtual wires, select the check box next to the virtual wire names and click **Delete**. Note that deleting a virtual wire removes it from the associated virtual wire interfaces shown on the **Interfaces** page.

## Packet Content Modification

Palo Alto Networks firewalls can be configured to perform the following packet content modifications:

- By default, the firewall will clear TCP urgent flags from all packets. This behavior can be changed by disabling this feature in the CLI using the following command: `set deviceconfig setting tcp urgent-data oobinline`.
- Modification of the DiffServ Code Point (DSCP) or IP precedences QoS markings can be configured in the **Actions** section of each security rule.
- VLAN tag translation will be performed whenever the firewall's input and output interfaces are tagged differently.

  ***Note:*** *VLAN tag translation is not available on virtual wire interfaces.*

# Firewall Interfaces

The following table describes the types of interfaces supported on the firewall and how to define them.

**Table 48. Supported Interfaces**

| Interface | Description |
|---|---|
| Layer 2 | One or more Layer 2 interfaces can be configured for untagged VLAN traffic. You can then define Layer 2 subinterfaces for traffic with specific VLAN tags. Refer to "Configuring Layer 2 Interfaces" on page 86 and "Configuring Layer 2 Subinterfaces" on page 87. |
| Layer 3 | One or more Layer 3 interfaces can be configured for untagged routed traffic. You can then define Layer 3 subinterfaces for traffic with specific VLAN tags.<br>Each interface can have multiple IP addresses. Refer to "Configuring Layer 3 Interfaces" on page 88 and "Configuring Layer 3 Subinterfaces" on page 91. |
| Aggregate Ethernet | Two or more Ethernet ports can be combined into a group to increase the throughput and resiliency for a Layer 2, Layer 3, or virtual wire interface and its subinterfaces. Refer to "Configuring Aggregate Ethernet Interfaces" on page 142.<br>***Note:*** *You cannot apply QoS settings to an aggregate Ethernet interface.* |
| VLAN | VLAN interfaces provide Layer 3 routing of VLAN traffic to non-VLAN destinations. Refer to "Configuring VLAN Interfaces" on page 143. |
| Loopback | Loopback interfaces can be used to provide Layer 3 services such as in-band management, GlobalProtect portal or gateway functionality, and IPSec. Each loopback interface behaves as a host interface and is assigned an IP address. Refer to "Configuring Loopback Interfaces" on page 146. |
| Tunnel | Tunnel interfaces can be configured. Refer to "Configuring Tunnel Interfaces" on page 147. |
| Virtual Wire | A virtual wire binds two Ethernet ports together, which allows you to install the firewall transparently in the network with minimum configuration. A virtual wire accepts all traffic or traffic with selected VLAN tags, but provides no switching or routing services. You can also create subinterfaces and classify traffic according to an IP address, IP range, or subnet. Refer to "Configuring Virtual Wire Interfaces" on page 138. |
| Tap | The Tap interface permits connection to a span port on a switch for traffic monitoring only. This mode does not support traffic blocking or URL filtering. Refer to "Configuring Tap Interfaces" on page 149. |
| High Availability | You can configure a data interface to be a high availability (HA) interface on some Palo Alto Networks firewalls. Refer to "Configuring HA Interfaces" on page 150. |

## Viewing the Current Interfaces

▶ *Network > Interfaces*

The **Interfaces** page lists the interface type, link state, and security zone for each configured interface, along with the IP address, virtual router, VLAN tag, and VLAN or virtual wire name (as applicable).

By default, the interfaces are listed by interface name.

The following icon is used on the **Interfaces** page:

Indicates the link is up (green), down (red), or in an unknown state (gray).

## Configuring Layer 2 Interfaces

▶ *Network > Interfaces > Ethernet*

You can configure one or more Ethernet ports as Layer 2 interfaces for untagged VLAN traffic. For each main Layer 2 interface, you can define multiple Layer 2 subinterfaces for traffic with specific VLAN tags (refer to "Configuring Layer 2 Subinterfaces" on page 129). VLAN interfaces can also be used in Layer 2 deployments to provide routing or gateway services to traffic in the Layer 2 domain (refer to "Configuring VLAN Interfaces" on page 143). This can be achieved by creating a VLAN interface and assigning it as the default gateway for hosts connected to the firewall's Layer2 interface.

To configure a Layer 2 Ethernet interface, click the link for the interface on the **Ethernet** tab, and specify the following settings.

**Table 49. Layer 2 Interface Settings**

| Field | Description |
|---|---|
| Interface Name | Choose the interface from the drop-down list. Modify the name if desired. |
| Interface Type | Select **Layer 2** from the drop-down list. |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Enter an optional description of the interface. |
| **Config Tab** | |
| VLAN | Select a VLAN, or click **New** to define a new VLAN (refer to "VLAN Support" on page 152). **None** removes the configuration from the interface. A VLAN object must be configured to enable switching between Layer 2 interfaces or to enable routing through a VLAN interface. |
| Virtual System | Select the virtual system for the interface. **None** removes the configuration from the interface. |
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). **None** removes the configuration from the interface. |
| **Advanced Tab** | |
| Link Speed | Select the interface speed in Mbps (10, 100, or 1000) or select auto. |

**Table 49. Layer 2 Interface Settings (Continued)**

| Field | Description |
|---|---|
| Link Duplex | Select whether the interface transmission mode is full-duplex (full), half-duplex (half), or negotiated automatically (auto). |
| Link State | Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto). |

## Configuring Layer 2 Subinterfaces

▶ *Network > Interfaces*

For each Ethernet port configured as a Layer 2 interface, you can define an additional logical Layer 2 interface (subinterface) for each VLAN tag that is used on the traffic received by the port. To configure the main Layer 2 interfaces, refer to "Configuring Layer 2 Interfaces" on page 128. To enable switching between Layer 2 subinterfaces, simply link those subinterfaces to the same VLAN object.

To add a Layer 2 Ethernet subinterface, select the associated physical interface and then click **Add Subinterface** and specify the following information.

**Table 50. Layer 2 Subinterface Settings**

| Field | Description |
|---|---|
| Interface Name | Select the Layer 2 interface where you want to add a subinterface. To configure the Layer 2 interfaces, refer to "Configuring Layer 2 Interfaces" on page 128.<br>Enter the number (1 to 9999) appended to the physical interface name to form the logical interface name. The general name format is:<br>ethernet*x*/*y*.<1-9999> |
| Tag | Enter the tag number (1 to 4094) of the traffic received on this interface. Outgoing traffic on this interface is also set to this tag value. |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Enter an optional description of the interface. |
| **Assign Interface To** | |
| VLAN | For a Layer 2 interface, select a VLAN, or click **New** to define a new VLAN (refer to "Network Profiles" on page 174). **None** removes the configuration from the interface. A VLAN object must be configured to enable switching between Layer 2 interfaces or to enable routing through a VLAN interface. |
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). **None** removes the configuration from the interface. |
| Virtual System | Select the virtual system for the interface. **None** removes the configuration from the interface. |

## Configuring Layer 3 Interfaces

▶ *Network > Interfaces > Ethernet*

You can configure one or more Ethernet ports as Layer 3 interfaces for untagged routed traffic. You can then define Layer 3 subinterfaces for traffic with specific VLAN tags (refer to "Configuring Layer 3 Subinterfaces" on page 134) when connecting the firewall to a neighboring device using a trunk link. For information on configuring Layer 3 interfaces for PPPoE, refer to "Point-to-Point Protocol over Ethernet Support" on page 124.
Layer 3 interfaces can also act as a DHCP client to receive its configuration from an external DHCP server.

To configure a Layer 3 Ethernet interface, click the link for the interface on the **Ethernet** tab, and specify the following settings.

**Table 51. Layer 3 Interface Settings**

| Field | Description |
|---|---|
| Interface Name | Choose the interface from the drop-down list. Modify the name if desired. |
| Interface Type | Select **Layer 3** from the drop-down list. |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Enter an optional description of the interface. |
| **Config Tab** | |
| Virtual Router | Select a virtual router, or click **New** to define a new virtual router (refer to "Virtual Routers and Routing Protocols" on page 153). **None** removes the configuration from the interface. |
| Virtual System | Select the virtual system for the interface. **None** removes the configuration from the interface. |
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). **None** removes the configuration from the interface. |
| **IPv4 Tab** | |
| Type | Choose how the IP address information will be specified (Static, PPPoE, or DHCP Client), as described below. |
| Static | Enter an IP address and network mask for the interface in the format *ip_address/mask*, and click **Add**. You can enter multiple IP addresses for the interface. To delete an IP address, select the address and click **Delete**. |

**Table 51. Layer 3 Interface Settings (Continued)**

| Field | Description |
|---|---|
| PPPoE | Choose **PPPoE** if the interface will be used for PPPoE and configure the following settings:<br>**General subtab:**<br>• **Enable**—Select the check box to activate the interface for PPPoE termination.<br>• **Username**—Enter the user name for the point-to-point connection.<br>• **Password/Confirm Password**—Enter and then confirm the password for the user name.<br>**Advanced subtab:**<br>• **Authentication**—Choose **CHAP** (Challenge-Handshake Authentication Protocol), **PAP** (Password Authentication Protocol), or the default **Auto** (to have the firewall determine the authentication protocol for PPPoE communications).<br>• **Static Address**—Specify the static IP address that was assigned by the service provider (optional, no default).<br>• **Automatically create default route pointing to peer**—Select the check box to automatically create a default route that points to the PPPoE peer when connected.<br>• **Default Route Metric**—Specify the route metric to be associated with the default route and used for path selection (optional, range 1-65535).<br>• **Access Concentrator**—Specify the name of the access concentrator to which the connection is made (optional, no default).<br>• **Service**—Specify the service string (optional, no default).<br>• **Passive**—Select the check box to use passive mode. In passive mode, a PPPoE end point waits for the access concentrator to send the first frame.<br><br>***Note:*** *PPPoE is not supported in HA active/active mode.* |
| DHCP Client | Choose **DHCP Client** to allow the interface to act as a DHCP client and receive a dynamically assigned IP address. Specify the following:<br>• **Enable**—Select the check box to activate the DHCP client on the interface.<br>• **Automatically create default route point to server**—Select the check box to automatically create a default route that points to the default gateway provided by the DHCP server.<br>• **Default Route Metric**—Specify the route metric to be associated with the default route and used for path selection (optional, range 1-65535).<br>Click **Show DHCP Client Runtime Info** to open a window that displays all settings received from the DHCP server, including DHCP lease status, dynamic IP assignment, subnet mask, gateway, server settings (DNS, NTP, domain, WINS, NIS, POP3, and SMTP).<br><br>***Note:*** *DHCP client is not supported in HA active/active mode.* |
| **IPv6 Tab** | |
| Enable IPv6 on the interface | Select the check box to enable IPv6 addressing on this interface. |

**Table 51. Layer 3 Interface Settings (Continued)**

| Field | Description |
|---|---|
| Interface ID | Enter the 64-bit extended unique identifier in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. If the interface ID is left blank, the firewall will use the EUI-64 generated from the physical interface's MAC address. The interface ID is used as the host portion of an interface address when the **Use interface ID as host portion** option is enabled. |
| Address | Click **Add** and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select **Anycast** to include routing through the nearest node. If the **Prefix** is not entered, the IPv6 address assigned to the interface will be wholly specified in the address text box.<br>Use the **Send Router Advertisement (Send RA)** option to enable router advertisement for this IP address. You can also set the **Autonomous** flag to be sent and you can set the **on-link** option. You must enable the global **Enable Router Advertisement** option on the interface before enabling **Send Router Advertisement** option for a specific IP address. |
| Address Resolution<br>(Duplicate Address Detection) | Select the check box to enable Duplicate Address Detection (DAD) and specify the following information.<br>• **DAD Attempts**—Specify the number of attempts within the neighbor solicitation interval for DAD before the attempt to identify neighbors fails (range 1-10).<br>• **Reachable Time**—Specify the length of time that a neighbor remains reachable after a successful query and response (range 1-36000 seconds).<br>• **Neighbor Solicitation (NS) Interval**—Specify the number of seconds for DAD attempts before failure is indicated (range 1-10 seconds). |

**Table 51. Layer 3 Interface Settings (Continued)**

| Field | Description |
|---|---|
| Enable Router Advertisement | Select the check box to enable Router Advertisement (RA) to provide Stateless Address Autoconfiguration (SLAAC) on IPv6 interfaces. This enables the firewall to act as a default gateway for IPv6 hosts that are not statically configured and will provide the host with an IPv6 prefix that can be used for address configuration. A separate DHCPv6 server can be used in conjunction with this feature to provide DNS and other settings to clients.<br>This option is a global setting for the interface, you can also set router advertisement options per IP address by clicking **Add** and entering in an IP address. You must enable this option on the interface if you are going to specify the **Send Router Advertisement** option per address.<br>Specify the following information that will be used by clients who receive the RA messages.<br>• **Min Interval (sec)**—Specify the minimum interval per second that the firewall will send out router advertisements. Router advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 3-1350 seconds, default 200 seconds).<br>• **Max Interval (sec)**—Specify the maximum interval per second that the firewall will send out router advertisements. Router advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 4-1800 seconds, default 600 seconds).<br>• **Hop Limit**—Specify the hop limit that will be applied to clients for outgoing packets. Enter 0 for no hop limit (range 1-255, default 64).<br>• **Link MTU**—Specify the link MTU that will be applied to clients. Select **unspecified** for no link MTU (range 1280-9192, default unspecified).<br>• **Reachable Time (ms)**—Specify the reachable time that the client will use to assume a neighbor is reachable after having received a reachability confirmation message. Select **unspecified** for no reachable time value (range 0-3600000 milliseconds, default unspecified).<br>• **Retrans Time (ms)**—Specify the retransmission timer that the client will use to determine how long it should wait before retransmitting neighbor solicitation messages. Select **unspecified** for no retrans time (range 0-4294967295 milliseconds, default unspecified).<br>• **Router Lifetime (sec)**—Specify the router lifetime that instructs the client on how long the firewall/router should be used as the default router (range 0-9000 seconds, default 1800).<br>• **Managed Configuration**—Select the check box to indicate to the client that addresses are available via DHCPv6.<br>• **Other Configuration**—Select the check box to indicate to the client that other addresses information is available via DHCPv6, such as DNS-related settings.<br>• **Consistency check**—Select the check box to enable consistency checks that the firewall will use to verify that router advertisement sent from other routers are advertising consistent information on the link. If inconsistencies are detected, a log will be created. |
| **Advanced Tab** | |
| Link Speed | Select the interface speed in Mbps (10, 100, or 1000) or select auto. |
| Link Duplex | Select whether the interface transmission mode is full-duplex (full), half-duplex (half), or negotiated automatically (auto). |

**Table 51. Layer 3 Interface Settings (Continued)**

| Field | Description |
|---|---|
| Link State | Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto). |
| Other Info | Specify the following information on the **Other Info** subtab:<br>• **Management Profile**—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.<br>• **MTU**—Enter the maximum transmission unit (MTU) in bytes for packets sent on this Layer 3 interface (512 to 1500, default 1500). If machines on either side of the firewall perform Path MTU Discovery (PMTUD), the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.<br>• **Adjust TCP MSS**—If you select this check box, the maximum segment size (MSS) is adjusted to 40 bytes less than the interface MTU. This setting addresses the situation in which a tunnel through the network requires a smaller MSS. If a packet cannot fit within the MSS without fragmenting, this setting allows an adjustment to be made.<br>• **Untagged Subinterface**—Specifies that all subinterfaces belonging to this Layer 3 interface are untagged.<br>PAN-OS selects an untagged subinterface as the ingress interface based on the destination of the packet. If a packet has an untagged subinterface's IP address as the destination, it will map to the subinterface. This also means that packets going in the reverse direction must have their source address translated to the untagged subinterface's interface IP address. A by product of this classification mechanism is that all multicast and broadcast packets will be assigned to the base interface rather than any of the subinterfaces. Since OSPF uses multicast, it is not supported on untagged subinterfaces. |
| ARP/Interface Entries | To add one or more static ARP entries, click **Add** and enter an IP address and its associated hardware (MAC) address and Layer 3 interface that can access the hardware address. |
| ND Entries | Click **Add** to enter the IPv6 address and MAC address of neighbors to add for discovery. |

## Configuring Layer 3 Subinterfaces

▶ *Network > Interfaces*

For each Ethernet port configured as a Layer 3 interface, you can define an additional logical Layer 3 interface (subinterface) for each VLAN tag that is used on the traffic received by the port. To configure the main Layer 3 interfaces, refer to "Configuring Layer 3 Interfaces" on page 130.

Untagged layer 3 subinterfaces may also be used when the parent Layer 3 interface's "untagged subinterface" option is enabled. Untagged subinterfaces are used in multi-tenant environments where each tenant's traffic must leave the firewall without VLAN tags.

Consider an example where each tenant's traffic egresses the firewall and the next hop is an ISP router. It is not always possible to apply a VLAN tag on the return traffic for proper classification into a virtual system by the firewall. In these cases, you can use an untagged subinterface on the ISP-router facing side. Each untagged subinterface will have an IP address and all outgoing traffic must have its source address translated to that interface IP address. An explicit NAT rule must be created for this feature to function. Source NAT is required on the untagged subinterfaces because the firewall will use the destination IP address on inbound

(return path) packets to select the appropriate virtual system for policy lookup. Any traffic received on the parent interface that is not destined for one of the untagged subinterface IPs will be handled by the virtual system and virtual router assigned to that parent interface.

To add a Layer 3 Ethernet subinterface, select the associated interface and click **Add Subinterface** and specify the following information.

**Table 52. Layer 3 Subinterface Settings**

| Field | Description |
|---|---|
| Interface Name | Select the Layer 3 interface where you want to add a subinterface. To configure the Layer 3 interfaces, refer to "Configuring Layer 3 Interfaces" on page 130.<br>Enter the number (1 to 9999) appended to the physical interface name to form the logical interface name. The general name format is:<br>ethernet*x*/*y*.<1-9999> |
| Tag | Enter the tag number (1 to 4094) of the traffic received on this interface. Outgoing traffic on this interface is also set to this tag value. |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Enter an optional description of the interface. |
| **Config Tab** | |
| Virtual Router | Select a virtual router, or click **New** to define a new virtual router (refer to "Virtual Routers and Routing Protocols" on page 153). **None** removes the configuration from the interface. |
| Virtual System | Select the virtual system for the interface. **None** removes the configuration from the interface. |
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). **None** removes the configuration from the interface. |
| **IPv4 Tab** | |
| Type | Choose how the IP address information will be specified (Static, PPPoE, or DHCP Client), as described below. |
| Static | Enter an IP address and network mask for the interface in the format *ip_address/mask*, and click **Add**. You can enter multiple IP addresses for the interface. To delete an IP address, select the address and click **Delete**. |
| DHCP Client | Choose **DHCP Client** to allow the interface to act as a DHCP client and receive a dynamically assigned IP address. Specify the following:<br>• **Enable**—Select the check box to activate the DHCP client on the interface.<br>• **Automatically create default route pointing to default gateway provided by server**—Select the check box to automatically create a default route that points to the DHCP server when connected.<br>• **Default Route Metric**—Specify the route metric to be associated with the default route and used for path selection (optional, range 1-65535).<br>• **Show DHCP Client Runtime Info**—Select to open a window that displays all settings received from the DHCP server, including DHCP lease status, dynamic IP assignment, subnet mask, gateway, server settings (DNS, NTP, domain, WINS, NIS, POP3, and SMTP). |

**Table 52. Layer 3 Subinterface Settings (Continued)**

| Field | Description |
|---|---|
| **IPv6 Tab** | |
| Enable IPv6 on the interface | Select the check box to enable IPv6 addressing on this interface. |
| Interface ID | Enter the 64-bit extended unique identifier in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. If the interface ID is left blank, the firewall will use the EUI-64 generated from the physical interface's MAC address. |
| Address | Click **Add** and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select **Anycast** to include routing through the nearest node. If **Prefix** is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.<br>Use the **Send Router Advertisement (Send RA)** option to enable router advertisement for this IP address. You can also set the **Autonomous** flag to be sent and you can set the **on-link** option. You must enable the global **Enable Router Advertisement** option on the interface before enabling **Send RA** option for a specific IP address. |
| Address Resolution<br>(Duplicate Address Detection) | Select the check box to enable Duplicate Address Detection (DAD) and specify the following information.<br>• **DAD Attempts**—Specify the number of attempts within the neighbor solicitation interval for DAD before the attempt to identify neighbors fails (range 1-10).<br>• **Reachable Time**—Specify the length of time that a neighbor remains reachable after a successful query and response (range 1-36000 seconds).<br>• **Neighbor Solicitation (NS) Interval**—Specify the number of seconds for DAD attempts before failure is indicated (range 1-10 seconds). |

**Table 52. Layer 3 Subinterface Settings (Continued)**

| Field | Description |
|---|---|
| Enable Router Advertisement | Select the check box to enable Router Advertisement (RA) to provide Stateless Address Autoconfiguration (SLAAC) on IPv6 interfaces. This enables the firewall to act as a default gateway for IPv6 hosts that are not statically configured and will provide the host with an IPv6 prefix that can be used for address configuration. A separate DHCPv6 server can be used in conjunction with this feature to provide DNS and other settings to clients.<br>This option is a global setting for the interface, you can also set router advertisement options per IP address by clicking **Add** and entering in an IP address. You must enable this option on the interface if you are going to specify the **Send Router Advertisement** option per address.<br>Specify the following information that will be used by clients who receive the RA messages.<br>• **Min Interval (sec)**—Specify the minimum interval per second that the firewall will send out router advertisements. Router Advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 3-1350 seconds, default 200 seconds).<br>• **Max Interval (sec)**—Specify the maximum interval per second that the firewall will send out router advertisements. Router Advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 4-1800 seconds, default 600 seconds).<br>• **Hop Limit**—Specify the hop limit that will be applied to clients for outgoing packets. Enter 0 for no hop limit (range 1-255, default 64).<br>• **Link MTU**—Specify the link MTU that will be applied to clients. Select **unspecified** for no link MTU (range 1280-9192, default unspecified).<br>• **Reachable Time (ms)**—Specify the reachable time that the client will use to assume a neighbor is reachable after having received a reachability confirmation message. Select **unspecified** for no reachable time value (range 0-3600000 milliseconds, default unspecified).<br>• **Retrans Time (ms)**—Specify the retransmission timer that the client will use to determine how long it should wait before retransmitting neighbor solicitation messages. Select **unspecified** for no retrans time (range 0-4294967295 milliseconds, default unspecified).<br>• **Router Lifetime (sec)**—Specify the router lifetime that instructs the client on how long the firewall/router should be used as the default router (range 0-9000 seconds, default 1800).<br>• **Managed Configuration**—Select the check box to indicate to the client that addresses are available via DHCPv6.<br>• **Other Configuration**—Select the check box to indicate to the client that other addresses information is available via DHCPv6, such as DNS-related settings.<br>• **Consistency check**—Select the check box to enable consistency checks that the firewall will use to verify that router advertisement sent from other routers are advertising consistent information on the link. If inconsistencies are detected, a log will be created. |

**Table 52. Layer 3 Subinterface Settings (Continued)**

| Field | Description |
|---|---|
| **Advanced Tab** | |
| Other Info | Specify the following information on the **Other Info** subtab:<br>• **Management Profile**—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.<br>• **MTU**—Enter the maximum transmission unit (MTU) in bytes for packets sent on this Layer 3 interface (512 to 1500, default 1500). If machines on either side of the firewall perform Path MTU Discovery (PMTUD), the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.<br>• **Adjust TCP MSS**—If you select this check box, the maximum segment size (MSS) is adjusted to 40 bytes less than the interface MTU. This setting addresses the situation in which a tunnel through the network requires a smaller MSS. If a packet cannot fit within the MSS without fragmenting, this setting allows an adjustment to be made. |
| ARP Entries | To add one or more static Address Resolution Protocol (ARP) entries, enter an IP address and its associated hardware (Media Access Control or MAC) address, and click **Add**. To delete a static entry, select the entry and click **Delete**. Static ARP entries reduce ARP processing and preclude man-in-the-middle attacks for the specified addresses. |
| ND Entries | Click **Add** to enter the IP address and MAC address of neighbors to add for discovery. |

## Configuring Virtual Wire Interfaces

▶ *Network > Interfaces*

To create a virtual wire, you bind two Ethernet ports together, which allows all traffic to pass between the ports, or just traffic with selected VLAN tags (no other switching or routing services are available). You can also create subinterfaces and classify traffic according to an IP address, IP range, or subnet. A virtual wire requires no changes to adjacent network devices. For an overview of virtual wire deployments, refer to "Virtual Wire Deployments" on page 120.

To set up a virtual wire through the firewall, you must first define the virtual wire interfaces, as described in the following procedure. You then create the virtual wire using the interfaces that you created.

To configure each virtual wire interface, follow these steps:

1. Identify the interface you want to use for the virtual wire on the **Ethernet** tab, and remove it from the current security zone, if any.
2. Click the interface name and specify the following information.

**Table 53. Virtual Wire Settings**

| Field | Description |
|---|---|
| Interface Name | The interface name is automatically populated based on the interface that you selected. This field cannot be edited. |
| Interface Type | Select **Virtual Wire** from the drop-down list. |

**Table 53. Virtual Wire Settings (Continued)**

| Field | Description |
|---|---|
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Enter an optional description of the interface. |
| **Config Tab** | |
| Virtual Wire | Select a virtual wire, or click **New** to define a new virtual wire (refer to "Defining Virtual Wires" on page 125). |
| Virtual System (only on systems with multi-virtual system capability) | Select the virtual system for the interface. |
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). |
| **Advanced Tab** | |
| Link Speed | Specify the interface speed. If the selected interface is a 10 Gbps interface, the only option is **auto**. In other cases, the options are: **10**, **100**, **1000**, or **auto**. |
| Link Duplex | Select whether the interface transmission mode is full-duplex (full), half-duplex (half), or negotiated automatically (auto). |
| Link State | Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto). |

If you want to change a virtual wire to another interface type, click the virtual wire name shown in the **VLAN/Virtual Wire** column, if any, select **None**, and click **OK**.

## Configuring Virtual Wire Subinterfaces

▶ *Network > Interfaces*

Using virtual wire subinterfaces allows you to separate traffic by VLAN tags or a VLAN tag and IP classifier combination, assign the tagged traffic to a different zone and virtual system, and then enforce security policies for the traffic that matches the defined criteria.

For each Ethernet port configured as a virtual wire interface, you can create one or more logical subinterfaces. Each subinterface can segregate traffic using a VLAN tag (with a zero or non-zero value, where the value zero indicates untagged traffic). After you set up a VLAN tagged subinterface, you can then create another subinterface that segregates traffic by combining the VLAN tag with IP classifiers such as a single IP address, a range, or a subnet.

IP classification on subinterfaces can only be used on one side of a virtual wire. Traffic entering the firewall through a virtual wire subinterface where IP classification is defined will be classified by its source IP address. For return path traffic entering through the other end of the virtual wire, classification is based on the destination IP address.

Make sure that the VLAN tags defined on the **Tag Allowed** list of the parent virtual wire interface (Network > Virtual Wires) are not included on the list of tags defined on a virtual wire subinterface.

To configure a parent or main virtual wire interfaces, refer to "Configuring Virtual Wire Interfaces" on page 138.

***Note:*** *The following NAT limitations apply to virtual wires:*

- Source NAT where the ingress interface uses source IP based classification is not supported.
- Destination NAT cannot be used where the egress interface uses source IP based classification.

To add a virtual wire subinterface, select the virtual wire interface where you want to add the subinterface, and click **Add Subinterface** and specify the following information.

**Table 54. Virtual Wire Subinterface Settings**

| Field | Description |
|---|---|
| Interface Name | The main interface name is automatically populated based on the interface that you selected; the label cannot be edited.<br>To define the subinterface, enter a number (1 to 9999) to the physical interface name to form the logical interface name. The general name format is:<br>ethernet*x*/*y*.<1-9999><br>To configure the virtual wire, refer to "Configuring Virtual Wire Interfaces" on page 138. |
| Tag | Enter the tag number (0 to 4094) of the traffic received on this interface.<br>Setting a tag value of 0 will match untagged traffic. |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Enter an optional description of the interface. |
| IP Classifier | Click **Add** to add an IP address, subnet, or IP range or any combination of the IP classifiers, to classify traffic entering the firewall through this physical port into this subinterface based on its source IP address. Return-path traffic entering the firewall through the other end of the associated virtual wire will be matched according to its destination address.<br>On a virtual wire subinterface, IP classification can only be used in conjunction with VLAN based classification. |
| **Assign Interface To** | |
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). |
| Virtual System (only on systems with multi-virtual system capability) | Select the virtual system for the interface. |
| Virtual Wire | Select a virtual wire, or click **New** to define a new virtual wire (refer to "Defining Virtual Wires" on page 125). |

## Configuring Aggregate Interface Groups

▶ *Network > Interfaces*

Aggregate interface groups allow you to generate more than 1 Gbps aggregate throughput by using 802.3ad link aggregation of multiple 1 Gbps links. Aggregation of 10Gbps XFP and SFP+ is also supported. The aggregate interface that you create becomes a logical interface. Interface management, zone profiles, VPN interfaces, and VLAN subinterfaces are all properties of the logical aggregate interface, not of the underlying physical interfaces.

Each aggregate group can contain several physical interfaces of the type Aggregate Ethernet. After the group is created, you perform operations such as configuring Layer 2 or Layer 3 parameters on the Aggregate Group object rather than on the Aggregate Ethernet interfaces themselves.

***Note:*** *The algorithm used to evenly distribute traffic on the interfaces in an aggregate group is based on the session ID of the incoming traffic. As traffic enters the aggregate group, the last three unique bits of the session ID are used to determine which interface should be used for a given flow.*

The following rules apply to aggregate interface groups:

- The interfaces are compatible with virtual wire, Layer 2, and Layer 3 interfaces.
- Tap mode is not supported.
- The 1 Gbps links in a group must be of the same type (all copper or all fiber).
- You can include up to eight aggregate interfaces in an aggregate group.
- All of the members of an aggregate group must be of the same type. This is validated during the commit operation.
- Aggregate groups can be used for redundancy and throughput scaling on the HA3 (packet forwarding) link in Active/Active HA deployments.

You can configure one or more interfaces as part of an aggregate Ethernet interface group. First define the group, as described in this section, and then assign interfaces to the group. For instructions on assigning interfaces to the group, refer to "Configuring Layer 3 Subinterfaces" on page 134.

To create and configure aggregate group interfaces, click **Add Aggregate Group** and specify the following information.

**Table 55. Aggregate Group Interface Settings**

| Field | Description |
|---|---|
| Interface Name | Enter a name and numeric suffix to identify the interface. The interface name is listed as *mm.n* where *mm* is the name and *n* is the suffix (1-8). |
| Interface Type | Select the interface type.<br>• **HA**—No additional configuration is required.<br>• **Layer 2**—Configure the settings as described in Table 50.<br>• **Layer 3**—Configure the settings as described in Table 52. |
| Comment | Enter an optional description of the interface. |

**Table 55. Aggregate Group Interface Settings (Continued)**

| Field | Description |
|---|---|
| **Assign Interface To** | |
| Assign Interface To | The interface assignment depends on the interface type, as follows:<br>• **Layer 2**—Specify a VLAN and zone.<br>• **Layer 3**—Specify a virtual router and zone<br>• **Virtual Wire**—Specify a virtual wire and zone.<br>***Note:*** *If the type is HA, there are no options to specify in this section.* |
| Virtual System | Select the virtual system for the interface. **None** removes the configuration from the interface. |

## Configuring Aggregate Ethernet Interfaces

▶ *Network > Interfaces*

Each aggregate Ethernet interface is assigned a name of the form ae.*number* and can be of the type Layer 2, Layer 3, or virtual wire. After the assignment is made, the new interface functions in the same way as any other interface.

To configure aggregate Ethernet interfaces, click the interface name on the **Ethernet** tab and specify the following information.

**Table 56. Aggregate Ethernet Interface Settings**

| Field | Description |
|---|---|
| Interface Name | Choose the interface from the drop-down list. Modify the name if desired. |
| Interface Type | Select **Aggregate Ethernet** from the drop-down list. |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Enter an optional description of the interface. |
| **Config Tab** | |
| Virtual System | Select the virtual system for the interface. **None** removes the configuration from the interface. |
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). **None** removes the configuration from the interface. |
| **Advanced Tab** | |
| Link Speed | Select the interface speed in Mbps (10, 100, or 1000) or auto. |
| Link Duplex | Select whether the interface transmission mode is full-duplex (full), half-duplex (half), or negotiated automatically (auto). |
| Link State | Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto). |

## Configuring VLAN Interfaces

▶ *Network > Interfaces*

For each Ethernet port configured as a Layer 2 interface, you can define a VLAN interface to allow routing of the VLAN traffic to Layer 3 destinations outside the VLAN. To configure the main Layer 2 interfaces, refer to "Configuring Layer 2 Interfaces" on page 128.

To define a VLAN interface, open the **VLAN** tab, click **Add**, and specify the following settings.

**Table 57. VLAN Interface Settings**

| Field | Description |
|---|---|
| Interface Name | Specify a numeric suffix for the interface (1-4999). |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Add an optional description of the interface. |
| **Config Tab** | |
| VLAN | Select a VLAN, or click **New** to define a new VLAN (refer to "Network Profiles" on page 174). **None** removes the configuration from the interface. |
| Virtual Router | Select a virtual router, or click **New** to define a new virtual router (refer to "Virtual Routers and Routing Protocols" on page 153). **None** removes the configuration from the interface. |
| Virtual System | Select the virtual system for the interface. **None** removes the configuration from the interface. |
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). **None** removes the configuration from the interface. |
| **IPv4 Tab** | |
| Static | Select **Static** to assign static IP addresses. Click **Add** and enter an IP address and network mask for the interface in the format *ip_address/mask*. You can enter multiple IP addresses for the interface. |
| DHCP Client | Select **DHCP** to use DHCP address assignment for the interface, and specify the following:<br>• **Enable**—Select the check box to activate the DHCP client on the interface.<br>• **Automatically create default route point to server**—Select the check box to automatically create a default route that points to the DHCP server when connected.<br>• **Default Route Metric**—Specify the route metric to be associated with the default route and used for path selection (optional, range 1-65535).<br>Click **Show DHCP Client Runtime Info** to open a window that displays all settings received from the DHCP server, including DHCP lease status, dynamic IP assignment, subnet mask, gateway, server settings (DNS, NTP, domain, WINS, NIS, POP3, and SMTP). |

**Table 57. VLAN Interface Settings (Continued)**

| Field | Description |
|---|---|
| ARP Entries | To add one or more static ARP entries, enter an IP address and its associated hardware (MAC) address, and click **Add**. To delete a static entry, select the entry and click **Delete**. |
| **IPv6 Tab** | |
| Enable IPv6 on the interface | Select the check box to enable IPv6 addressing for the subinterface. |
| Interface ID | Enter the 64-bit extended unique identifier in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. If the interface ID is left blank, the firewall will use the EUI-64 generated from the physical interface's MAC address. |
| Address | Click **Add** and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select **Anycast** to include routing through the nearest node. If **Prefix** is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.<br>Use the **Send Router Advertisement (Send RA)** option to enable router advertisement for this IP address. You can also set the **Autonomous** flag to be sent and you can set the **on-link** option. You must enable the global **Enable Router Advertisement** option on the interface before enabling **Send Router Advertisement** option for a specific IP address. |
| Address Resolution (Duplicate Address Detection | Select the check box to enable Duplicate Address Detection (DAD) and specify the following information.<br>• **DAD Attempts**—Specify the number of attempts within the neighbor solicitation interval for DAD before the attempt to identify neighbors fails (range 1-10).<br>• **Reachable Time**—Specify the length of time that a neighbor remains reachable after a successful query and response (range 1-36000 seconds).<br>**Neighbor Solicitation (NS) Interval**—Specify the number of seconds for DAD attempts before failure is indicated (range 1-10 seconds). |

**Table 57. VLAN Interface Settings (Continued)**

| Field | Description |
|---|---|
| Enable Router Advertisement | Select the check box to enable Router Advertisement (RA) to provide Stateless Address Autoconfiguration (SLAAC) on IPv6 interfaces. This enables the firewall to act as a default gateway for IPv6 hosts that are not statically configured and will provide the host with an IPv6 prefix that can be used for address configuration. A separate DHCPv6 server can be used in conjunction with this feature to provide DNS and other settings to clients.<br>This option is a global setting for the interface, you can also set router advertisement options per IP address by clicking **Add** and entering in an IP address. You must enable this option on the interface if you are going to specify the **Send Router Advertisement** option per address.<br>Specify the following information that will be used by clients who receive the RA messages.<br>• **Min Interval (sec)**—Specify the minimum interval per second that the firewall will send out router advertisements. Router Advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 3-1350 seconds, default 200 seconds).<br>• **Max Interval (sec)**—Specify the maximum interval per second that the firewall will send out router advertisements. Router Advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 4-1800 seconds, default 600 seconds).<br>• **Hop Limit**—Specify the hop limit that will be applied to clients for outgoing packets. Enter 0 for no hop limit (range 1-255, default 64).<br>• **Link MTU**—Specify the link MTU that will be applied to clients. Select **unspecified** for no link MTU (range 1280-9192, default unspecified).<br>• **Reachable Time (ms)**—Specify the reachable time that the client will use to assume a neighbor is reachable after having received a reachability confirmation message. Select **unspecified** for no reachable time value (range 0-3600000 milliseconds, default unspecified).<br>• **Retrans Time (ms)**—Specify the retransmission timer that the client will use to determine how long it should wait before retransmitting neighbor solicitation messages. Select **unspecified** for no retrans time (range 0-4294967295 milliseconds, default unspecified).<br>• **Router Lifetime (sec)**—Specify the router lifetime that instructs the client on how long the firewall/router should be used as the default router (range 0-9000 seconds, default 1800).<br>• **Managed Configuration**—Select the check box to indicate to the client that addresses are available via DHCPv6.<br>• **Other Configuration**—Select the check box to indicate to the client that other addresses information is available via DHCPv6, such as DNS-related settings.<br>• **Consistency check**—Select the check box to enable consistency checks that the firewall will use to verify that router advertisement sent from other routers are advertising consistent information on the link. If inconsistencies are detected, a log will be created. |

**Table 57. VLAN Interface Settings (Continued)**

| Field | Description |
|---|---|
| **Advanced Tab** | |
| Other Info | Specify the following:<br>• **Management Profile**—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.<br>• **MTU**—Enter the MTU in bytes for packets sent on this interface (512-1500, default 1500). If machines on either side of the firewall perform PMTUD, the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.<br>• **Adjust TCP MSS**—if you select this check box, the maximum segment size (MSS) is adjusted to 40 bytes less than the interface MTU. This setting addresses the situation in which a tunnel through the network requires a smaller MSS. If a packet cannot fit within the MSS without fragmenting, this setting allows an adjustment to be made. |
| ARP/Interface Entries | To add one or more static ARP entries, click **Add** and enter an IP address and its associated hardware (MAC) address and Layer 3 interface that can access the hardware address. |
| ND Entries | Click **Add** to enter the IP address and MAC address of neighbors to add for discovery. |

## Configuring Loopback Interfaces

▶ *Network > Interfaces*

You can define one or more Layer 3 loopback interfaces, as needed. For example, you can define a loopback interface to manage the firewall instead of using the management port.

To define a loopback interface, open the **Loopback** tab, click **Add**, and specify the following settings.

**Table 58. Loopback Interface Settings**

| Field | Description |
|---|---|
| Interface Name | Specify a numeric suffix for the interface (1-4999). |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Add an optional description of the interface. |
| **Config Tab** | |
| Virtual Router | Select a virtual router, or click **New** to define a new virtual router (refer to "Virtual Routers and Routing Protocols" on page 153). **None** removes the configuration from the interface. |
| Virtual System | Select the virtual system for the interface. **None** removes the configuration from the interface. |
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). **None** removes the configuration from the interface. |

**Table 58. Loopback Interface Settings (Continued)**

| Field | Description |
|---|---|
| **IPv4 Tab** | |
| IP Address | Click **Add** to enter IP addresses and network masks for the interface. |
| **IPv6 Tab** | |
| Enable IPv6 on the interface | Select the check box to enable IPv6 addressing for the subinterface. |
| Interface ID | Specify the unique 64-bit hexadecimal identifier for the subinterface. |
| Address | Enter the IPv6 address. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select **Anycast** to include routing through the nearest node. |
| **Advanced Tab** | |
| Other Info | Specify the following settings:<br>• **Management Profile**—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.<br>• **MTU**—Enter the maximum transmission unit (MTU) in bytes for packets sent on this interface (512 to 1500, default 1500). If machines on either side of the firewall perform Path MTU Discovery (PMTUD), the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.<br>• **Adjust TCP MSS**—If you select this check box, the maximum segment size (MSS) is adjusted to 40 bytes less than the interface MTU. This setting addresses the situation in which a tunnel through the network requires a smaller MSS. If a packet cannot fit within the MSS without fragmenting, this setting allows an adjustment to be made. |

## Configuring Tunnel Interfaces

▶ *Network > Interfaces*

To define tunnel interfaces, open the **Tunnel** tab, click **Add**, and specify the following settings.

**Table 59. Tunnel Interface Settings**

| Field | Description |
|---|---|
| Interface Name | Specify a numeric suffix for the interface (1-4999). |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Add an optional description of the interface. |
| **Config Tab** | |
| Virtual Router | Select a virtual router for this interface, or click **New** to configure a new virtual router. Refer to "Virtual Routers and Routing Protocols" on page 153. **None** removes the configuration from the interface. |
| Virtual System | Select the virtual system for the interface. **None** removes the configuration from the interface. |

**Table 59. Tunnel Interface Settings (Continued)**

| Field | Description |
|---|---|
| Security Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). **None** removes the configuration from the interface. |
| **IPv4 Tab** | |
| IP Address | Click **Add** to enter IP addresses and network masks for the interface. |
| **IPv6 Tab** | |
| Enable IPv6 on the interface | Select the check box to enable IPv6 addressing for the interface.<br>This option allows you to route IPv6 traffic over an IPv4 IPSec tunnel and will provide confidentiality between IPv6 networks. The IPv6 traffic is encapsulated by IPv4 and then ESP.<br>To route IPv6 traffic to the tunnel, you will either use a static route to the tunnel, or use a Policy Based Forwarding (PBF) rule to direct traffic and to provide redundancy by monitoring the other end of the tunnel and failing over when needed. |
| Interface ID | Enter the 64-bit extended unique identifier in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. If the interface ID is left blank, the firewall will use the EUI-64 generated from the physical interface's MAC address. |
| Address | Click **Add** and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select **Anycast** to include routing through the nearest node. If **Prefix** is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box. |
| **Advanced Tab** | |
| Other Info | Specify the following:<br>• **Management Profile**—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.<br>• **MTU**—Enter the MTU in bytes for packets sent on this interface (512-1500, default 1500). If machines on either side of the firewall perform PMTUD, the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.<br>***Note:*** *The firewall automatically considers tunnel overhead when performing IP fragmentation and also adjusts the TCP maximum segment size (MSS) as needed.* |

## Configuring Tap Interfaces

▶ *Network > Interfaces*

You can define tap interfaces as needed to permit connection to a span port on a switch for traffic monitoring only (refer to "Tap Mode Deployments" on page 125).

To define tap interfaces, click an interface name on the **Ethernet** tab, and specify the following information.

**Table 60. Tap Interface Settings**

| Field | Description |
|---|---|
| Interface Name | Specify a name for the interface or keep the default name. |
| Interface Type | Select **Tap** from the drop-down list. |
| Netflow Profile | Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85. |
| Comment | Enter an optional description of the interface. |
| **Config Tab** | |
| Virtual System | Select a virtual system. **None** removes the configuration from the interface. |
| Zone | Select a security zone for the interface, or click **New** to define a new zone (refer to "Defining Security Zones" on page 151). **None** removes the configuration from the interface. |
| **Advanced Tab** | |
| Link Speed | Select the interface speed in Mbps (10, 100, or 1000) or auto. |
| Link Duplex | Select whether the interface transmission mode is full-duplex (full), half-duplex (half), or negotiated automatically (auto). |
| Link State | Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto). |

## Configuring HA Interfaces

Each HA interface has a specific function: one interface is for configuration synchronization and heartbeats and the other interface is for state synchronization. If active/active high availability is enabled, a third HA interface can be used to forward packets.

***Note:*** *Some Palo Alto Networks firewalls include dedicated physical ports for use in HA deployments (one for the control link and one for the data link). For firewalls that do not include dedicated ports, you must specify the data ports that will be used for HA. For additional information on HA, refer to "Enabling HA on the Firewall" on page 101.*

To define HA interfaces, click an interface name and specify the following information.

**Table 61. HA Interface Settings**

| Field | Description |
|---|---|
| Interface Name | Choose the interface from the drop-down list. Modify the name if desired. |
| Interface Type | Select **HA** from the drop-down list. |
| Comment | Enter an optional description of the interface. |
| **Advanced Tab** | |
| Link Speed | Select the interface speed in Mbps (10, 100, or 1000) or auto. |
| Link Duplex | Select whether the interface transmission mode is full-duplex (full), half-duplex (half), or negotiated automatically (auto). |
| Link State | Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto). |

# Security Zones

A security zone identifies one or more source or destination interfaces on the firewall. When you define a security policy rule, you must specify the source and destination security zones of the traffic. For example, an interface connected to the Internet is in an "untrusted" security zone, while an interface connected to the internal network is in a "trusted" security zone.

Separate zones must be created for each type of interface (Layer 2, Layer 3, tap, or virtual wire), and each interface must be assigned to a zone before it can process traffic. Security policies can be defined only between zones of the same type. However, if you create a VLAN interface for one or more VLANs, applying security policies between the VLAN interface zone and a Layer 3 interface zone (Figure 16) has the same effect as applying policies between the Layer 2 and Layer 3 interface zones.

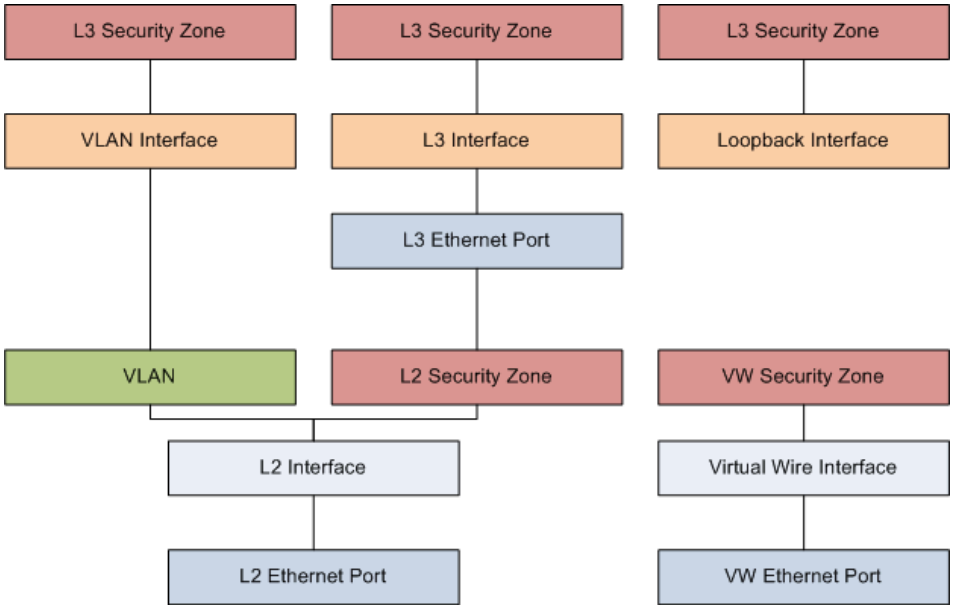


**Figure 16. Zone and Interface Types**

## Defining Security Zones

▶ *Network* > *Zones*

In order for a firewall interface to be able to process traffic, it must be assigned to a security zone. To define security zones, click **New** and specify the following information:

**Table 62. Security Zone Settings**

| Field | Description |
|---|---|
| Name | Enter a zone name (up to 15 characters). This name appears in the list of zones when defining security policies and configuring interfaces. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, periods, and underscores. |
| Location | Select the virtual system that applies to this zone. |

**Table 62. Security Zone Settings (Continued)**

| Field | Description |
|---|---|
| Type | Select a zone type (Layer2, Layer3, Virtual Wire, Tap, or External virtual system) to list all the interfaces of that type that have not been assigned to a zone. The Layer 2 and Layer 3 zone types list all Ethernet interfaces and subinterfaces of that type. The External virtual system type is for communications among virtual systems in the firewall. Refer to "Communications Among Virtual Systems" on page 111.<br>Each interface can belong to one zone in one virtual system. |
| Zone Protection Profiles | Select a profile that specifies how the security gateway responds to attacks from this zone. To add new profiles, refer to "Defining Zone Protection Profiles" on page 178. |
| Log Setting | Select a log forwarding profile for forwarding zone protection logs to an external system. |
| Enable User Identification | Select to enable the user identification function on a per-zone basis. |
| User Identification ACL Include List | Enter the IP address or IP address/mask of a user or group to be identified (format *ip_address/mask*; for example, 10.1.1.1/24). Click **Add**. Repeat as needed. If an include list is not configured, then all IP addresses are allowed. |
| User Identification ACL Exclude List | Enter the IP address or IP address/mask of a user or group that will explicitly not be identified (format *ip_address/mask*; for example, 10.1.1.1/24). Click **Add**. Repeat as needed. If an exclude list is not configured, then all IP addresses are allowed. |

# VLAN Support

▶ ***Network > VLANs***

The firewall supports VLANs that conform to the IEEE 802.1Q standard. Each Layer 2 interface that is defined on the firewall must be associated with a VLAN. The same VLAN can be assigned to multiple Layer 2 interfaces, but each interface can belong to only one VLAN. Optionally, a VLAN can also specify a VLAN interface that can route traffic to Layer 3 destinations outside the VLAN.

**Table 63. VLAN Settings**

| Field | Description |
|---|---|
| Name | Enter a VLAN name (up to 31 characters). This name appears in the list of VLANs when configuring interfaces. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| VLAN Interface | Select a VLAN interface to allow traffic to be routed outside the VLAN. To define a VLAN interface, refer to "Configuring VLAN Interfaces" on page 143. |
| L3 Forwarding Enabled | If you select a VLAN interface, you can select the check box to enable Layer 3 routing over the selected interface. |

**Table 63. VLAN Settings (Continued)**

| Field | Description |
|---|---|
| Interfaces | Specify firewall interfaces for the VLAN. |
| Static MAC Configuration | Specify the interface through which a MAC address is reachable. This will override any learned interface-to-MAC mappings. |

# Virtual Routers and Routing Protocols

You can set up virtual routers to enable the firewall to route packets at Layer 3 by making packet forwarding decisions according to the destination IP address. The Ethernet interfaces and VLAN interfaces defined on the firewall receive and forward the Layer 3 traffic. The destination zone is derived from the outgoing interface based on the forwarding criteria, and policy rules are consulted to identify the security policies to be applied. In addition to routing to other network devices, virtual routers can route to other virtual routers within the same firewall if a next hop is specified to point to another virtual router.

Support is provided for static routing and dynamic routing using the Routing Information Protocol (RIP), Open Shortest Path First (OSPF) protocol, and Border Gateway Protocol (BGP).

***Note:*** *Policy-based forwarding is also supported for traffic on Layer 3 interfaces.*

## Routing Information Protocol

RIP was designed for small IP networks and relies on hop count to determine routes; the best routes have the fewest number of hops. RIP is based on UDP and uses port 520 for route updates. By limiting routes to a maximum of 15 hops, the protocol helps prevent the development of routing loops, but also limits the supported network size. If more than 15 hops are required, traffic is not routed. RIP also can take longer to converge than OSPF and other routing protocols. The firewall supports RIP v2.

## Open Shortest Path First

OSPF determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The router keeps information about the links between it and the destination and can make highly efficient routing decisions. A cost is assigned to each router interface, and the best routes are determined to be those with the lowest costs, when summed over all the encountered outbound router interfaces and the interface receiving the LSA.

Hierarchical techniques are used to limit the number of routes that must be advertised and the associated LSAs. Because OSPF dynamically processes a considerable amount of route information, it has greater processor and memory requirements than does RIP.

## Border Gateway Protocol

The Border Gateway Protocol (BGP) is the primary Internet routing protocol. BGP determines network reachability based on IP prefixes that are available within autonomous systems (AS), where an AS is a set of IP prefixes that a network provider has designated to be part of a single routing policy.

In the routing process, connections are established between BGP peers (or neighbors). If a route is permitted by the policy, it is stored in the routing information base (RIB). Each time the local firewall RIB is updated, the firewall determines the optimal routes and sends an update to the external RIB, if export is enabled.

Conditional advertisement is used to control how BGP routes are advertised. The BGP routes must satisfy conditional advertisement rules before being advertised to peers.

BGP supports the specification of aggregates, which combine multiple routes into a single route. During the aggregation process, the first step is to find the corresponding aggregation rule by performing a longest match that compares the incoming route with the prefix values for other aggregation rules.

The firewall provides a complete BGP implementation that includes the following features:

- Specification of one BGP routing instance per virtual router.
- Routing policies based on route-map to control import, export and advertisement, prefix-based filtering, and address aggregation.
- Advanced BGP features that include route reflector, AS confederation, route flap dampening, and graceful restart.
- IGP-BGP interaction to inject routes to BGP using redistribution profiles.

BGP configuration consists of the following elements:

- Per-routing-instance settings, which include basic parameters such as local route ID and local AS and advanced options such as path selection, route reflector, AS confederation, route flap, and dampening profiles.
- Authentication profiles, which specify the MD5 authentication key for BGP connections.
- Peer group and neighbor settings, which include neighbor address and remote AS and advanced options such as neighbor attributes and connections.
- Routing policy, which specifies rule sets that peer groups and peers use to implement imports, exports, conditional advertisements, and address aggregation controls.

## Multicast Routing

The multicast routing feature allows the firewall to route multicast streams using Protocol Independent Multicast Sparse Mode (PIM-SM) and PIM Source Specific Multicast (PIM-SSM) for applications such as media broadcasting (radio and video) with PIMv2. The firewall performs Internet Group Management Protocol (IGMP) queries for hosts that are on the same network as the interface on which IGMP is configured. PIM-SM and IGMP can be enabled on Layer 3 interfaces. IGMP v1, v2, and v3 are supported. PIM and IGMP must be enabled on host-facing interfaces.

PAN-OS provides full multicast security while acting as a PIM designated router (DR), PIM rendezvous point (RP), intermediate PIM router, or IGMP querier. The firewall can be deployed in environments in which the RP is statically configured or dynamically elected. The

bootstrap router (BSR) role is not supported. Deployment across IPSec tunnels is fully supported between Palo Alto Networks firewalls. GRE encapsulation within IPSec is not currently supported.

### Security policy

PAN-OS provides two methods to enforce security on multicast feeds. Multicast groups can be filtered in the IGMP and PIM group permission settings specified on an interface level. Multicast traffic must also be explicitly allowed by security policy. A special destination zone known as "Multicast" has been added and must be specified to control multicast traffic in security, QoS, and DoS protection rules. In contrast to unicast security policy, multicast security policies must be explicitly created when the source and destination interfaces are in the same zone. Security profiles are supported in multicast environments that require threat prevention capabilities.

### Logging

Each multicast session passing through the firewall creates only one traffic log entry (even if the firewall is replicating packets for distribution on multiple interfaces). Traffic logs indicate the number of bytes coming into the firewall rather than the number of bytes distributed as part of the multicast feed.

## Defining Virtual Routers

▶ *Network > Virtual Routers*

Defining virtual routers allows you to set up forwarding rules for Layer 3 and enable the use of dynamic routing protocols. Each Layer 3 interface, loopback interface, and VLAN interface defined on the firewall should be associated with a virtual router. Each interface can belong to only one virtual router.

***Note:*** *To configure Ethernet ports as Layer 3 interfaces, refer to "Configuring Layer 3 Interfaces" on page 130. To define Layer 3 subinterfaces, refer to "Configuring Layer 3 Subinterfaces" on page 134. For an overview of virtual routers, refer to "Virtual Routers and Routing Protocols" on page 153.*

Define settings on the specified tabs, as appropriate.

### General Tab

Select the interfaces to include in the virtual router and add any static routes. Refer to the following table.

**Table 64. Virtual Router Settings - General Tab**

| Field | Description |
|---|---|
| Name | Specify a name to describe the virtual router (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Interfaces | Select the interfaces that you want to include in the virtual router. When you select an interface, it is included in the virtual router and can be used as an outgoing interface in the virtual router's routing tab.<br>To specify the interface type, refer to "Firewall Interfaces" on page 127.<br>***Note:*** *When you add an interface, its connected routes are added automatically.* |

**Table 64. Virtual Router Settings - General Tab (Continued)**

| Field | Description |
|---|---|
| Administrative Distances | Specify the following administrative distances:<br>• Static routes (10-240, default 10).<br>• OSPF Int (10-240, default 30).<br>• OSPF Ext (10-240, default 110).<br>• IBGP (10-240, default 200).<br>• EBGP (10-240, default 20).<br>• RIP (10-240, default 120). |

### Static Routes Tab

Optionally enter one or more static routes. Click the IP or IPv6 tab to specify the route using IPv4 or IPv6 addresses. It is usually necessary to configure default routes (0.0.0.0/0) here. Default routes are applied for destinations that are otherwise not found in the virtual router's routing table.

**Table 65. Virtual Router Settings - Static Routes Tab**

| Field | Description |
|---|---|
| Name | Enter a name to identify the static route (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Destination | Enter an IP address and network mask in the format *ip_address/mask*. |
| Interface | Select the interface to forward packets to the destination, or configure the next hop settings, or both. |
| Next Hop | Specify the following next hop settings:<br>• **None**—Select if there is no next hop for the route.<br>• **IP Address**—Specify the IP address of the next hop router.<br>• **Discard**—Select if you want to drop traffic that is addressed to this destination.<br>• **Next VR**—Select a virtual router in the firewall as the next hop. This option allows you to route internally between virtual routers within a single firewall. |
| Admin Distance | Specify the administrative distance for the static route (10-240, default 10). |
| Metric | Specify a valid metric for the static route (1 - 65535). |
| No Install | Select if you do not want to install the route in the forwarding table. The route is retained in the configuration for future reference. |

### Redistribution Profiles Tab

Modify route redistribution filter, priority, and action based on desired network behavior. Route redistribution allows static routes and routes that are acquired by other protocols to be advertised through specified routing protocols. Redistribution profiles must be applied to routing protocols in order to take effect. Without redistribution rules, each protocol runs separately and does not communicate outside its purview. Redistribution profiles can be added or modified after all routing protocols are configured and the resulting network topology is established. Apply redistribution profiles to the RIP and OSPF protocols by defining export rules. Apply redistribution profiles to BGP in the Redistribution Rules tab. Refer to the following table.

**Table 66. Virtual Router Settings - Redistribution Profiles Tab**

| Field | Description |
|---|---|
| Name | Click **Add** to display the **Redistribution Profile** page, and enter the profile name. |
| Priority | Enter a priority (range 1-255) for this profile. Profiles are matched in order (lowest number first). |
| Redistribute | Choose whether to perform route redistribution based on the settings in this window.<br>• **Redist**—Select to redistribute matching candidate routes. If you select this option, enter a new metric value. A lower metric value means a more preferred route.<br>• **No Redist**—Select to not redistribute matching candidate routes. |
| **General Filter Tab** | |
| Type | Select check boxes to specify the route types of the candidate route. |
| Interface | Select the interfaces to specify the forwarding interfaces of the candidate route. |
| Destination | To specify the destination of the candidate route, enter the destination IP address or subnet (format x.x.x.x or x.x.x.x/n) and click **Add**. To remove an entry, click the icon associated with the entry. |
| Next Hop | To specify the gateway of the candidate route, enter the IP address or subnet (format x.x.x.x or x.x.x.x/n) that represents the next hop and click **Add**. To remove an entry, click the icon associated with the entry. |
| **OSPF Filter Tab** | |
| Path Type | Select check boxes to specify the route types of the candidate OSPF route. |
| Area | Specify the area identifier for the candidate OSPF route. Enter the OSPF area ID (format x.x.x.x), and click **Add**. To remove an entry, click the icon associated with the entry. |
| Tag | Specify OSPF tag values. Enter a numeric tag value (1-255), and click Add. To remove an entry, click the icon associated with the entry. |
| **BGP Filter Tab** | |
| Community | Specify a community for BGP routing policy. |
| Extended Community | Specify an extended community for BGP routing policy. |

### RIP Tab

Specify parameters for use of the Routing Information Protocol (RIP) on the selected interfaces. Although it is possible to configure both RIP and OSPF, it is generally recommended to choose only one of these protocols. Refer to the following table.

**Table 67. Virtual Router Settings - RIP Tab**

| Field | Description |
|---|---|
| Enable | Select the check box to enable the RIP protocol. |
| Reject Default Route | Select the check box if you do not want to learn any default routes through RIP. Selecting the check box is highly recommended. |
| Allow Redist Default Route | Select the check box to permit redistribution of default routes through RIP. |
| **Interfaces** | |
| Interface | Select the interface that runs the RIP protocol. |
| Enable | Select to enable these settings. |
| Advertise | Select to advertise a default route to RIP peers with the specified metric value. |
| Metric | Specify a metric value for the router advertisement. This field is visible only if the **Advertise** check box is selected. |
| Auth Profile | Select the profile. |
| Mode | Select **normal**, **passive**, or **send-only**. |
| **Timers** | |
| Interval Seconds (sec) | Define the length of the timer interval in seconds. This duration is used for the remaining RIP timing fields (1 - 60). |
| Update Intervals | Enter the number of intervals between route update announcements (1 - 3600). |
| Expire Intervals | Enter the number of intervals between the time that the route was last updated to its expiration (1- 3600). |
| Delete Intervals | Enter the number of intervals between the time that the route expires to its deletion (1- 3600). |
| **Auth Profiles** | |
| Profile Name | Enter a name for the authentication profile to authenticate RIP messages. To authenticate RIP messages, first define the authentication profiles and then apply them to interfaces on the **RIP** tab. |
| Password Type | Select the type of password (simple or MD5).<br>• If you select **Simple**, enter the simple password and then confirm.<br>• If you select **MD5**, enter one or more password entries, including **Key-ID** (0-255), **Key**, and optional **Preferred** status. Click **Add** for each entry, and then click **OK**. To specify the key to be used to authenticate outgoing message, select the **Preferred** option. |

**Table 67. Virtual Router Settings - RIP Tab (Continued)**

| Field | Description |
|---|---|
| **Export Rules** | |
| Export Rules | (Read-only) Displays the rules that apply to routes sent by the virtual router to a receiving router.<br>• **Allow Redistribute Default Route**—Select the check box to permit the firewall to redistribute its default route to peers.<br>• **Redistribution Profile**—Select a redistribution profile that allows you to modify route redistribution, filter, priority, and action based on the desired network behavior. Refer to "Redistribution Profiles Tab" on page 158. |

## OSPF Tab

Specify parameters for use of the Open Shortest Path First (OSPFv2) protocol on the selected interfaces. Although it is possible to configure both RIP and OSPF, it is generally recommended to choose only one IGP protocol. Refer to the following table.BGP Tab

**Table 68. Virtual Router Settings - OSPF Tab**

| Field | Description |
|---|---|
| Enable | Select the check box to enable the OSPF protocol. |
| Reject Default Route | Select the check box if you do not want to learn any default routes through OSPF. Selecting the check box is recommended, especially for static routes. |
| Router ID | Specify the router ID associated with the OSPF instance in this virtual router. The OSPF protocol uses the router ID to uniquely identify the OSPF instance. |
| **Areas** | |
| Area ID | Configure the area over which the OSPF parameters can be applied.<br>Enter an identifier for the area in x.x.x.x format. This is the identifier that each neighbor must accept to be part of the same area. |

**Table 68. Virtual Router Settings - OSPF Tab (Continued)**

| Field | Description |
|---|---|
| Type | Select one of the following options.<br>• **Normal**—There are no restrictions; the area can carry all types of routes.<br>• **Stub**—There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, select **Accept Summary** if you want to accept this type of link state advertisement (LSA) from other areas. Also, specify whether to include a default route LSA in advertisements to the stub area along with the associated metric value (1-255). If the **Accept Summary** option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.<br>• **NSSA** (Not-So-Stubby Area)—It is possible to leave the area directly, but only by routes other than OSPF routes. If you select this option, select **Accept Summary** if you want to accept this type of LSA. Specify whether to include a default route LSA in advertisements to the stub area along with the associated metric value (1-255). Also, select the route type used to advertise the default LSA. Click **Add** in the **External Ranges** section and enter ranges if you want to enable or suppress advertising external routes that are learned through NSSA to other areas. |
| Range | Click **Add** to aggregate LSA destination addresses in the area into subnets. Enable or suppress advertising LSAs that match the subnet, and click **OK**. Repeat to add additional ranges. |
| Interface | Click **Add** and enter the following information for each interface to be included in the area, and click **OK**.<br>• **Interface**—Choose the interface.<br>• **Enable**—Cause the OSPF interface settings to take effect.<br>• **Passive**—Select the check box to if you do not want the OSPF interface to send or receive OSPF packets. Although OSPF packets are not sent or received if you choose this option, the interface is included in the LSA database.<br>• **Link type**—Choose **Broadcast** if you want all neighbors that are accessible through the interface to be discovered automatically by multicasting OSPF hello messages, such as an Ethernet interface. Choose **p2p** (point-to-point) to automatically discover the neighbor. Choose **p2mp** (point-to-multipoint) when neighbors must be defined manually. Defining neighbors manually is allowed only for p2mp mode.<br>• **Metric**—Enter the OSPF metric for this interface (0-65535).<br>• **Priority**—Enter the OSPF priority for this interface (0-255). It is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR) according to the OSPF protocol. When the value is zero, the router will not be elected as a DR or BDR.<br>• **Auth Profile**—Select a previously-defined authentication profile.<br>• **Timing**—It is recommended that you keep the default timing settings.<br>• **Neighbors**—For p2pmp interfaces, enter the neighbor IP address for all neighbors that are reachable through this interface. |

**Table 68. Virtual Router Settings - OSPF Tab (Continued)**

| Field | Description |
|---|---|
| Virtual Link | Configure the virtual link settings to maintain or enhance backbone area connectivity. The settings must be defined for area boarder routers, and must be defined within the backbone area (0.0.0.0). Click **Add**, enter the following information for each virtual link to be included in the backbone area, and click **OK**.<br>• **Name**—Enter a name for the virtual link.<br>• **Neighbor ID**—Enter the router ID of the router (neighbor) on the other side of the virtual link.<br>• **Transit Area**—Enter the area ID of the transit area that physically contains the virtual link.<br>• **Enable**—Select to enable the virtual link.<br>• **Timing**—It is recommended that you keep the default timing settings.<br>• **Auth Profile**—Select a previously-defined authentication profile. |
| **Auth Profiles** | |
| Profile Name | Enter a name for the authentication profile. To authenticate the OSPF messages, first define the authentication profiles and then apply them to interfaces on the **OSPF** tab. |
| Password Type | Select the type of password (simple or MD5).<br>• If you select **Simple**, enter the password.<br>• If you select **MD5**, enter one or more password entries, including **Key-ID** (0-255), **Key**, and optional **Preferred** status. Click **Add** for each entry, and then click **OK**. To specify the key to be used to authenticate outgoing message, select the **Preferred** option. |
| **Export Rules** | |
| Allow Redistribute Default Route | Select the check box to permit redistribution of default routes through OSPF. |
| Name | Select the name of a redistribution profile. The value must be an IP subnet or valid redistribution profile name. |
| New Path Type | Choose the metric type to apply. |
| New Tag | Specify a tag for the matched route that has a 32-bit value. |
| **Advanced** | |
| RFC 1583 Compatibility | Select the check box to assure compatibility with RFC 1583. |
| Timers | • **SPF Calculation Delay (sec)**—This option is a delay timer allowing you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times.<br>• **LSA Interval (sec)**—The option specifies the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur. |

### BGP Tab

Specify parameters for use of Border Gateway Protocol (BGP) on the selected interfaces. Refer to the following table.

**Table 69. Virtual Router Settings - BGP Tab**

| Field | Description |
|---|---|
| Enable | Select the check box to enable BGP. |
| Router ID | Enter the IP address to assign to the virtual router. |
| AS Number | Enter the number of the AS to which the virtual router belongs, based on the router ID (range 1-4294967295). |
| **General Tab** | |
| Reject Default Route | Select the check box to ignore any default routes that are advertised by BGP peers. |
| Install Route | Select the check box to install BGP routes in the global routing table. |
| Aggregate MED | Select to enable route aggregation even when routes have different Multi-Exit Discriminator (MED) values. |
| Default Local Preference | Specifies a value than can be used to determine preferences among different paths. |
| AS Format | Select the 2-byte (default) or 4-byte format. This setting is configurable for interoperability purposes. |
| Always Compare MED | Enable MED comparison for paths from neighbors in different autonomous systems. |
| Deterministic MED Comparison | Enable MED comparison to choose between routes that are advertised by IBGP peers (BGP peers in the same autonomous system). |
| Auth Profiles | Click **Add** to include a new authentication profile and configure the following settings:<br>• **Profile Name**—Enter a name to identify the profile.<br>• **Secret/Confirm Secret**—Enter and confirm a passphrase for BGP peer communications.<br>Click the icon to delete a profile. |
| **Advanced Tab** | |
| Graceful Restart | Activate the graceful restart option.<br>• **Stale Route Time**—Specify the length of time that a route can stay in the stale state (range 1-3600 seconds, default 120 seconds).<br>• **Local Restart Time**—Specify the length of time that the local device takes to restart. This value is advertised to peers (range 1-3600 seconds, default 120 seconds).<br>• **Max Peer Restart Time**—Specify the maximum length of time that the local device accepts as a grace period restart time for peer devices (range 1-3600 seconds, default 120 seconds). |
| Reflector Cluster ID | Specify an IPv4 identifier to represent the reflector cluster. |
| Confederation Member AS | Specify the identifier for the AS confederation to be presented as a single AS to external BGP peers. |

**Table 69. Virtual Router Settings - BGP Tab (Continued)**

| Field | Description |
|---|---|
| Dampening Profiles | Settings include:<br>• **Profile Name**—Enter a name to identify the profile.<br>• **Enable**—Activate the profile.<br>• **Cutoff**—Specify a route withdrawal threshold above which a route advertisement is suppressed (range 0.0-1000.0, default 1.25).<br>• **Reuse**—Specify a route withdrawal threshold below which a suppressed route is used again (range 0.0-1000.0, default 5).<br>• **Max. Hold Time**—Specify the maximum length of time that a route can be suppressed, regardless of how unstable it has been (range 0-3600 seconds, default 900 seconds).<br>• **Decay Half Life Reachable**—Specify the length of time after which a route's stability metric is halved if the route is considered reachable (range 0-3600 seconds, default 300 seconds).<br>• **Decay Half Life Unreachable**—Specify the length of time after which a route's stability metric is halved if the route is considered unreachable (range 0-3600 seconds, default 300 seconds).<br>Click the icon to delete a profile. |
| **Peer Group Tab** | |
| Name | Enter a name to identify the peer. |
| Enable | Select to activate the peer. |
| Aggregated Confed AS Path | Select the check box to include a path to the configured aggregated confederation AS. |
| Soft Reset with Stored Info | Select the check box to perform a soft reset of the firewall after updating the peer settings. |
| Type | Specify the type of peer or group and configure the associated settings (see below in this table for descriptions of Import Next Hop and Export Next Hop).<br>• **IBGP**—Specify the following;<br>– Export Next Hop<br>• **EBGP Confed**—Specify the following;<br>– Export Next Hop<br>• **IBGP Confed**—Specify the following;<br>– Export Next Hop<br>• **EBGP**—Specify the following:<br>– Import Next Hop<br>– Export Next Hop<br>– Remove Private AS (select if you want to force BGP to remove private AS numbers). |
| Import Next Hop | Choose an option for next hop import:<br>• **original**—Use the Next Hop address provided in the original route advertisement.<br>• **use-peer**—Use the peer's IP address as the Next Hop address. |

**Table 69. Virtual Router Settings - BGP Tab (Continued)**

| Field | Description |
|---|---|
| Export Next Hop | Choose an option for next hop export:<br>• **resolve**—Resolve the Next Hop address using the local forwarding table.<br>• **use-self**—Replace the Next Hop address with this router's IP address to ensure that it will be in the forwarding path. |
| Peer | To add a new peer, click **New** and configure the following settings:<br>• **Name**—Enter a name to identify the peer.<br>• **Enable**—Select to activate the peer.<br>• **Peer AS**—Specify the AS of the peer.<br>• **Local Address**—Choose a firewall interface and local IP address.<br>• **Connection Options**—Specify the following options:<br>– **Auth Profile**—Select the profile.<br>– **Keep Alive Interval**—Specify an interval after which routes from a peer are suppressed according to the hold time setting (range 0-1200 seconds, default 30 seconds).<br>– **Multi Hop**—Set the time-to-live (TTL) value in the IP header (range 1-255, default 0). The default value of 0 means 2 for eBGP and 255 for iBGP.<br>– **Open Delay Time**—Specify the delay time between opening the peer TCP connection and sending the first BGP open message (range 0-240 seconds, default 0 seconds).<br>– **Hold Time**—Specify the period of time that may elapse between successive KEEPALIVE or UPDATE messages from a peer before the peer connection is closed. (range 3-3600 seconds, default 90 seconds).<br>– **Idle Hold Time**—Specify the time to wait in the idle state before retrying connection to the peer (range 1-3600 seconds, default 15 seconds).<br>• **Peer Address**—Specify the IP address and port of the peer.<br>• **Advanced Options**—Configure the following settings:<br>– **Reflector Client**—Select the type of reflector client (**Non-Client**, **Client**, or **Meshed Client**). Routes that are received from reflector clients are shared with all internal and external BGP peers.<br>– **Peering Type**—Specify a bilateral peer, or leave unspecified.<br>– **Max. Prefixes**—Specify the maximum number of supported IP prefixes (1 - 100000 or unlimited).<br>• **Incoming Connections/Outgoing Connections**—Specify the incoming and outgoing port numbers and select the Allow check box to allow traffic to or from these ports. |

**Table 69. Virtual Router Settings - BGP Tab (Continued)**

| Field | Description |
|---|---|
| **Import Rules/Export Rules Tabs** | |
| Import Rules/Export Rules | Click the BGP **Import Rules** or **Export Rules** subtab. To add a new rule, click **Add** and configure the following settings.<br>• **General** subtab:<br>– **Name**—Specify a name to identify the rule.<br>– **Enable**—Select to activate the rule.<br>– Used **by**—Select the peer groups that will use this rule.<br>• **Match** subtab:<br>– **AS-Path Regular Expression**—Specify a regular expression for filtering of AS paths.<br>– **Community Regular Expression**—Specify a regular expression for filtering of community strings.<br>– **Extended Community Regular Expression**—Specify a regular expression for filtering of extended community strings.<br>– **Address Prefix**—Specify IP addresses or prefixes for route filtering.<br>– **MED**—Specify a MED value for route filtering.<br>– **Next Hop**—Specify next hop routers or subnets for route filtering.<br>– **From Peer**—Specify peer routers for route filtering.<br>• **Action** subtab:<br>– **Action**—Specify an action (**Allow** or **Deny**) to take when the match conditions are met.<br>– **Local Preference**—Specify a local preference metric, only if the action is **Allow**.<br>– **MED**—Specify a MED value, only if the action is **Allow** (0- 65535).<br>– **Weight**—Specify a weight value, only if the action is **Allow** (0- 65535).<br>– **Next Hop**—Specify a next hop router, only if the action is **Allow**.<br>– **Origin**—Specify the path type of the originating route: IGP, EGP, or incomplete, only if the action is **Allow**.<br>– **AS Path Limit**—Specify an AS path limit, only if the action is **Allow**.<br>– **AS Path**—Specify an AS path: **None**, **Remove**, **Prepend**, **Remove and Prepend**, only if the action is **Allow**.<br>– **Community**—Specify a community option: **None**, **Remove All**, **Remove Regex**, **Append**, or **Overwrite**, only if the action is **Allow**.<br>– **Extended Community**—Specify a community option: **None**, **Remove All**, **Remove Regex**, **Append**, or **Overwrite**, only if the action is **Allow**.<br>– **Dampening**—Specify the dampening parameter, only if the action is **Allow**.<br>Click the icon to delete a group. Click **Clone** to add a new group with the same settings as the selected group. A suffix is added to the new group name to distinguish it from the original group. |

**Table 69. Virtual Router Settings - BGP Tab (Continued)**

| Field | Description |
|---|---|
| **Conditional Adv Tab** | The BGP conditional advertisement feature allows you to control what route to advertise in the event that a different route is not available in the local BGP routing table (LocRIB), indicating a peering or reachability failure. This is useful in cases where you want to try and force routes to one AS over another, for example if you have links to the Internet through multiple ISPs and you want traffic to be routed to one provider instead of the other unless there is a loss of connectivity to the preferred provider. With conditional advertising, you can configure a non-exist filter that matches the prefix of the preferred route. If any route matching the non-exist filter is not found in the local BGP routing table, only then will the device allow advertisement of the alternate route (the route to the other, non-preferred provider) as specified in its advertise filter. To configure conditional advertisement, select the **Conditional Adv** tab and then click **Add**. The following describes how to configure the values in the fields. |
| **Policy** | Specify the policy name for this conditional advertisement rule. |
| **Enable** | Select the check box to enable BGP conditional advertisement. |
| **Used By** | Click **Add** and select the peer groups that will use this conditional advertisement policy. |
| **Non Exist Filters Subtab** | Use this tab to specify the prefix(es) of the preferred route. This specifies the route that you want to advertise, if it is available in the local BGP routing table. If a prefix is going to be advertised and matches a Non Exist filter, the advertisement will be suppressed.<br>Click **Add** to create a non-exist filter.<br>• **Non Exist Filters**—Specify a name to identify this filter.<br>• **Enable**—Select to activate the filter.<br>• **AS-Path Regular Expression**—Specify a regular expression for filtering of AS paths.<br>• **Community Regular Expression**—Specify a regular expression for filtering of community strings.<br>• **Extended Community Regular Expression**—Specify a regular expression for filtering of extended community strings.<br>• **MED**—Specify a MED value for route filtering.<br>• **Address Prefix**—Click **Add** and then specify the exact NLRI prefix for the preferred route.<br>• **Next Hop**—Specify next hop routers or subnets for route filtering.<br>• **From Peer**—Specify peer routers for route filtering. |

**Table 69. Virtual Router Settings - BGP Tab (Continued)**

| Field | Description |
|---|---|
| **Advertise Filters Subtab** | Use this tab to specify the prefix(es) of the route in the Local-RIB routing table that should be advertised in the event that the route in the non-exist filter is not available in the local routing table.<br>If a prefix is going to be advertised and does not match a Non Exist filter, the advertisement will occur.<br>Click **Add** to create an advertise filter.<br>• **Advertise Filters**—Specify a name to identify this filter.<br>• **Enable**—Select to activate the filter.<br>• **AS-Path Regular Expression**—Specify a regular expression for filtering of AS paths.<br>• **Community Regular Expression**—Specify a regular expression for filtering of community strings.<br>• **Extended Community Regular Expression**—Specify a regular expression for filtering of extended community strings.<br>• **MED**—Specify a MED value for route filtering.<br>• **Address Prefix**—Click **Add** and then specify the exact NLRI prefix for the route to be advertised if the preferred route is not available.<br>• **Next Hop**—Specify next hop routers or subnets for route filtering.<br>• **From Peer**—Specify peer routers for route filtering. |
| **Aggregate Tab** | |
| Aggregate Filters | To add a new rule, click **Add** to display the settings. Configure the settings in the **Suppress Filters**, **Advertise Filters**, and **Aggregate Route Attributes** subtabs, and click **Done** to add the rule to the Addresses list. The parameters are described above in this table for the **Import Rules** and **Export Rules** tabs.<br>Click the icon to delete a rule. |
| **Redist Rules Tab** | |
| Name | Select the name of a redistribution profile. The value must be an IP subnet or valid redistribution profile name. |
| Allow Redistribute Default Route | Select the check box to permit the firewall to redistribute its default route to BGP peers. |
| Redist Rules | To add a new rule, click **Add**, configure the settings, and click **Done**. The parameters are described above in this table for the **Import Rules** and **Export Rules** tabs.<br>Click the icon to delete a rule. |

### Multicast Tab

Specify settings for multicast routing in the following table.

**Table 70. Virtual Router Settings - Multicast Tab**

| Field | Description |
|---|---|
| Enable | Select the check box to enable multicast routing. |
| **Rendezvous Point Subtab** | |
| RP Type | Choose the type of Rendezvous Point (RP) that will run on this virtual router. A static RP must be explicitly configured on other PIM routers whereas a candidate RP is elected automatically.<br>• **None**—Choose if there is no RP running on this virtual router.<br>• **Static**—Specify a static IP address for the RP and choose options for **RP Interface** and **RP Address** from the drop-down lists. Select the **Override learned RP for the same group** check box if you want to use the specified RP instead of the RP elected for this group.<br>• **Candidate**—Specify the following information for the candidate RP running on this virtual router:<br>– **RP Interface**—Select an interface for the RP. Valid interface types include loopback, L3, VLAN, aggregate Ethernet, and tunnel.<br>– **RP Address**—Select an IP address for the RP.<br>– **Priority**—Specify a priority for candidate RP messages (default 192).<br>– **Advertisement interval**—Specify an interval between advertisements for candidate RP messages.<br>• **Group list**—If you choose **Static** or **Candidate**, click **Add** to specify a list of groups for which this candidate RP is proposing to be the RP. |
| Remote Rendezvous Point | Click **Add** and specify the following:<br>• **IP address**—Specify the IP address for the RP.<br>• **Override learned RP for the same group**—Select the check box to use the specified RP instead of the RP elected for this group.<br>• **Group**—Specify a list of groups for which the specified address will act as the RP. |
| **Interfaces Subtab** | |
| Name | Enter a name to identify an interface group. |
| Description | Enter an optional description. |
| Interface | Click **Add** to specify one or more firewall interfaces. |
| Group Permissions | Specify general rules for multicast traffic:<br>• **Any Source**—Click **Add** to specify a list of multicast groups for which PIM-SM traffic is permitted.<br>• **Source-Specific**—Click **Add** to specify a list of multicast group and multicast source pairs for which PIM-SSM traffic is permitted. |

**Table 70. Virtual Router Settings - Multicast Tab (Continued)**

| Field | Description |
|---|---|
| IGMP | Specify rules for IGMP traffic. IGMP must be enabled for host facing interfaces (IGMP router) or for IGMP proxy host interfaces.<br>• **Enable**—Select the check box to enable the IGMP configuration.<br>• **IGMP Version**—Choose version 1, 2, or 3 to run on the interface.<br>• **Enforce Router-Alert IP Option**—Select the check box to require the router-alert IP option when speaking IGMPv2 or IGMPv3. This option must be disabled for compatibility with IGMPv1.<br>• **Robustness**—Choose an integer value to account for packet loss on a network (range 1-7, default 2). If packet loss is common, choose a higher value.<br>• **Max Sources**—Specify the maximum number of source-specific memberships allowed on this interface (0 = unlimited).<br>• **Max Groups**—Specify the maximum number of groups allowed on this interface.<br>• **Query Configuration**—Specify the following:<br>– **Query interval**—Specify the interval at which general queries are sent to all hosts.<br>– **Max Query Response Time**—Specify the maximum time between a general query and a response from a host.<br>– **Last Member Query Interval**—Specify the interval between group or source-specific query messages (including those sent in response to leave-group messages).<br>– **Immediate Leave**—Select the check box to leave the group immediately when a leave message is received. |
| PIM configuration | Specify the following Protocol Independent Multicast (PIM) settings:<br>• **Enable**—Select the check box to allow this interface to receive and/or forward PIM messages<br>• **Assert Interval**—Specify the interval between PIM assert messages.<br>• **Hello Interval**—Specify the interval between PIM hello messages.<br>• **Join Prune Interval**—Specify the interval between PIM join and prune messages (seconds). Default is 60.<br>• **DR Priority**—Specify the designated router priority for this interface<br>• **BSR Border**—Select the check box to use the interface as the bootstrap border.<br>• **PIM Neighbors**—Click **Add** to specify the list of neighbors that will communicate with using PIM. |
| **SPT Threshold Subtab** | |
| Name | The Shortest Path Tree (SPT) threshold defines the throughput rate (in kbps) at which multicast routing will switch from shared tree distribution (sourced from the rendezvous point) to source tree distribution.<br>Click **Add** to specify the following SPT settings:<br>• **Multicast Group Prefix**—Specify the multicast IP address/prefix for which the SPT will be switched to source tree distribution when the throughput reaches the desired threshold (kbps).<br>• **Threshold**—Specify the throughput at which we'll switch from shared tree distribution to source tree distribution |

**Table 70. Virtual Router Settings - Multicast Tab (Continued)**

| Field | Description |
|---|---|
| **Source Specific Address Space Subtab** | |
| Name | Defines the multicast groups for which the firewall will provide source-specific multicast (SSM) services.<br>Click **Add** to specify the following settings for source-specific addresses:<br>• **Name**—Enter a name to identify this group of settings.<br>• **Group**—Specify groups for the SSM address space.<br>• **Included**—Select this check box to include the specified groups in the SSM address space. |

### Displaying Runtime Statistics for Virtual Routers

▶ *Network > Virtual Routers*

Detailed runtime statistics are available for the virtual router and dynamic routing protocols from the **Virtual Routers** page. In the **Runtime Stats** column, click the **More Runtime Stats** link to open a new window that contains the routing table as well as routing protocol-specific details. For an overview of virtual routers, refer to "Virtual Routers and Routing Protocols" on page 153.

# DHCP Server and Relay

▶ *Network > DHCP*

The firewall supports the selection of DHCP servers or DHCP relay for IP address assignment on the Layer 3 and VLAN interfaces. Multiple DHCP servers are supported. Client requests can be forwarded to all servers, with the first server response sent back to the client.

The DHCP assignment also works across an IPSec VPN, allowing clients to receive an IP address assignment from a DHCP server on the remote end of an IPSec tunnel. For information on IPSec VPN tunnels, refer to "Configuring IPSec Tunnels" on page 309.

The settings depend on whether you select **DHCP Server** or **DHCP Relay** as the type.

**Table 71. DHCP Settings**

| Field | Description |
|---|---|
| **DHCP Server Tab** | |
| Interface | Select the firewall interface. |
| Mode | Choose whether the settings on this page are enabled, disabled, or are determined automatically. |
| Ping IP when allocating new IP | Select the check box to send a ping message when allocating a new IP address. |

**Table 71. DHCP Settings (Continued)**

| Field | Description |
|---|---|
| Lease | Select **Unlimited**, or select **Timeout** and enter any limitations on the DHCP lease interval. You can enter days, hours, or minutes. For example, if you enter only hours, then the lease is restricted to that number of hours. |
| Inheritance Source | Select a source to propagate various server settings from a DHCP client interface or PPPoE client interface into the DHCP server. Once an inheritance source is specified, select the desired fields in the DHCP server configuration and set to **inherited** to inherit the values from the DHCP server. For example, you can inherit DNS, WINS, NIS, NTP. |
| Primary DNS<br>Secondary DNS | Enter the IP address of the preferred and alternate Domain Name System (DNS) servers. The alternate server address is optional. |
| Primary WINS<br>Secondary WINS | Enter the IP address of the preferred and alternate Windows Internet Naming Service (WINS) servers. The alternate server address is optional. |
| Primary NIS<br>Secondary NIS | Enter the IP address of the preferred and alternate Network Information Service (NIS) servers. The alternate server address is optional. |
| Primary NTP<br>Secondary NTP | Enter the IP address of the preferred and alternate Network Time Protocol server. The alternate server address is optional. |
| Gateway | Enter the IP address of the network gateway that is used to reach the DHCP servers. |
| POP3 Server | Enter the IP address of the Post Office Protocol (POP3) server. |
| SMTP Server | Enter the IP address of the Simple Mail Transfer Protocol (SMTP) server. |
| DNS Suffix | Enter a suffix for the client to use locally when an unqualified hostname is entered that it cannot resolve. |
| IP Pools | Specify the range of IP addresses to which this DHCP configuration applies and click **Add**. You can enter an IP subnet and subnet mask (for example, 192.168.1.0/24) or a range of IP addresses (for example, 192.168.1.10-192.168.1.20). Add multiple entries to specify multiple IP address pools.<br>To edit an existing entry, click **Edit**, make the changes, and click **Done**. To delete an entry, click **Delete**.<br>***Note:*** *If you leave this area blank, there will be no restrictions on the IP ranges.* |
| Reserved Address | Specify the IP address (format *x.x.x.x*) or MAC address (format *xx:xx:xx:xx:xx:xx*) of any devices that you do not want to subject to DHCP address assignment.<br>To edit an existing entry, click **Edit**, make the changes, and click **Done**. To delete an entry, click **Delete**.<br>***Note:*** *If you leave this area blank, then there will be no reserved IP addresses.* |
| **DHCP Relay Tab** | |
| Interface | Select the firewall interface. |
| IPv4 | Select the **IPv4** check box to use IPv4 addresses for DHCP relay and specify IPv4 addresses for up to four DHCP servers. |
| IPv6 | Select the **IPv6** check box to use IPv6 addresses for DHCP relay and specify IPv6 addresses for up to four DHCP servers. Specify an outgoing interface if you are using an IPv6 multicast address for your server. |

# DNS Proxy

▶ *Network > DNS Proxy*

For all DNS queries that are directed to an interface IP address, the firewall supports the selective directing of queries to different DNS servers based on full or partial domain names. TCP or UDP DNS queries are sent through the configured interface. UDP queries switch over to TCP when a DNS query answer is too long for a single UDP packet.

If the domain name is not found in the DNS proxy cache, the domain name is searched for a match based on configuration of the entries in the specific DNS proxy object (on the interface on which the DNS query arrived) and forwarded to a name server based on the match results. If no match is found, the default name servers are used. Static entries and caching are also supported.

**Table 72. DNS Proxy Settings**

| Field | Description |
|---|---|
| Name | Specify a name to identify the DNS proxy (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Enable | Select the check box to enable DNS proxy. |
| Inheritance Source | Select a source to inherit default DNS server settings. This is commonly used in branch office deployments where the firewall's WAN interface is addressed by DHCP or PPPoE. |
| Primary<br>Secondary | Specify the IP addresses of the default primary and secondary DNS servers. If the primary DNS server cannot be found, the secondary will be used. |
| Check inheritance source | Click the link to see the server settings that are currently assigned to the DHCP client and PPPoE client interfaces. These may include DNS, WINS, NTP, POP3, SMTP, or DNS suffix. |
| Interface | Select the **Interface** check box to specify the firewall interfaces to support the DNS proxy rules. Select an interface from the drop-down list and click **Add**. You can add multiple interfaces. To delete an interface, select the interface and click **Delete**. |
| DNS Proxy Rules | Identify DNS proxy server rules. Click **Add** and specify the following information:<br>• **Name**—A name is required so that a static entry can be reference and modified via the CLI.<br>• **Turn on caching of domains resolved by this mapping**—Select the check box to enable caching of domains that are resolved by this mapping.<br>• **Domain Name**—Click **Add** and enter the proxy server domain name. Repeat to add additional names. To delete a name, select the name and click **Delete**. For a DNS proxy rule, the number of tokens in a wildcard string must match the number of tokens in the requested domain. For example, "*.engineering.local" will not match "engineering.local". Both must be specified.<br>• **Primary/Secondary**—Enter the hostname or IP addresses of the primary and secondary DNS servers. |

**Table 72. DNS Proxy Settings (Continued)**

| Field | Description |
|---|---|
| Static Entries | Provide static FQDN to IP address mappings that will be delivered in response to DNS queries made by hosts. Click Add and specify the following information:<br>• **Name**—Enter a name for the Static Entry.<br>• **FQDN**—Enter the Fully Qualified Domain Name (FQDN) that will be mapped to the static IP addresses defined in the Address field.<br>• **Address**—Click Add and enter the IP addresses that map to this domain.<br>Repeat to add additional addresses. To delete an address, select the address and click Delete. |
| Advanced | Specify the following information:<br>• **Cache**—Select the check box to enable DNS caching and specify the following information:<br>– **Size**—Specify the number of entries that the cache will hold (range 1024-10240, default 1024).<br>– **Timeout**—Specify the length of time (hours) after which all cached entries are removed. DNS time-to-live values are used to remove cache entries when they have been stored for less than the configured timeout period. Following a timeout, new requests must be resolved and cached again (range 4 to 24, default 4 hours).<br>• **TCP Queries**—Select the check box to enable DNS queries using TCP and specify the following information:<br>– **Max Pending Requests**—Specify the upper limit on the number of concurrent pending TCP DNS requests that the firewall will support (range 64-256, default 64).<br>• **UDP Queries Retries**—Specify settings for UDP query retries:<br>– **Interval**—Specify the time in seconds after which another request is sent if no response has been received (range 1-30, default 2 seconds).<br>– **Attempts**—Specify the maximum number of attempts (excluding the first attempt) after which the next DNS server is tried (range 1-30, default 5). |

# Network Profiles

Network profiles capture configuration information that the firewall can use to establish network connections and implement policies. The following types of network profiles are supported:

- **IKE gateways**, **IPSec crypto profiles**, and **IKE crypto profiles**—These profiles support configuration and operation of IPSec VPNs. For information on the following profile types, refer to "Configuring IPSec Tunnels" on page 309.
  - IKE gateways include the configuration information that is necessary to perform IKE protocol negotiation with peer gateways when setting up IPSec VPN tunnels.
  - IKE crypto profiles specify the protocols and algorithms for Phase 1 identification, authentication, and encryption in VPN tunnels.
  - IPSec crypto profiles specify the protocols and algorithms for Phase 2 identification, authentication, and encryption in VPN tunnels.
- **Monitor profiles—**These profiles are used to monitor IPSec tunnels and to monitor a next-hop device for policy based forwarding (PBF) rules. In both cases, the monitor profile is used to specify an action to take when a resource (IPSec tunnel or next-hop device) becomes unavailable. Refer to "Monitor Settings" on page 177.
- **Interface management profiles**—These profiles specify the protocols that can be used to manage the firewall for Layer 3 interfaces, including VLAN and loopback interfaces. Refer to "Defining Interface Management Profiles" on page 175.
- **Zone protection profiles**—These profiles determine how the firewall responds to attacks from individual security zones. Refer to "Defining Zone Protection Profiles" on page 178. The following types of protection are supported:
  - **Flood Protection**—Protects against SYN, ICMP, UDP, and other IP-based flooding attacks.
  - **Reconnaissance detection**—Allows you to detect and block commonly used port scans and IP address sweeps that attackers run to find potential attack targets.
  - **Packet-based attack protection**—Protects against large ICMP packets and ICMP fragment attacks as well as a number of IP and TCP level attacks. Per-zone non-syn-tcp behavior can also be specified.
  - **QoS profiles**—These profiles determine how the QoS traffic classes are treated. You can set overall limits on bandwidth regardless of class and also set limits for individual classes. You can also assign priorities to different classes. Priorities determine how traffic is treated in the presence of contention. Refer to "Defining QoS Profiles" on page 358.

## Defining Interface Management Profiles

▶ ***Network > Network Profiles > Interface Mgmt***

Use this page to specify the protocols that are used to manage the firewall. To assign management profiles to each interface, refer to "Configuring Layer 3 Interfaces" on page 130 and "Configuring Layer 3 Subinterfaces" on page 134. For an overview of firewall interfaces, refer to "Firewall Interfaces" on page 127.

***Note:*** *Although the default port used for HTTPS access is 443, if you enable HTTPS access on an in-band port and Global Protect is enabled on the same port, the port used for HTTPS access is automatically set to 4443 instead of 443.*

**Table 73. Interface Management Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of interface management profiles when configuring interfaces. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Ping<br>Telnet<br>SSH<br>HTTP<br>HTTPS<br>SNMP<br>Response Pages<br>User-ID Service | Select the check box for each service to be enabled on the interfaces where the profile is applied.<br>The **Response Pages** check box controls whether the ports used to serve captive portal and URL filtering response pages are open on Layer 3 interfaces. Ports 6080 and 6081 are left open if this setting is enabled.<br>The **User-ID Service** option is needed to allow communication between firewalls when a firewall is acting as a redistribution point to provide user mapping information to other PAN-OS firewalls. Refer to "User-ID Agent" on page 283. |
| Permitted IP Addresses | Enter the list of IPv4 or IPv6 addresses from which firewall management is allowed. |

## Defining Monitor Profiles

▶ *Network > Network Profiles > Monitor*

A monitor profile is used to monitor IPSec tunnels and to monitor a next-hop device for policy-based forwarding (PBF) rules. In both cases, the monitor profile is used to specify an action to take when a resource (IPSec tunnel or next-hop device) becomes unavailable. Monitor profiles are optional, but can be very useful for maintaining connectivity between sites and to ensure that PBF rules are maintained.

After creating a monitor profile, you can apply it to an IPSec Tunnel profile from the **General** tab, then selecting the **Show Advanced Options** check box, click the **Tunnel Monitor** check box and then select the desired profile from the drop-down list. In a PBF rule, click the **Forwarding** tab, click the **Monitor** check box and then select the desired profile from the drop down list.

In the example of an IPSec tunnel, the firewall monitors the specified IP address through the tunnel to determine if the tunnel is working properly according to the defined monitor profile. If the tunnel monitor IP address is not reachable, the action will be taken based on the settings you choose that are defined in "Monitor Settings" on page 177.

For a monitor profiles used in PBF rule, a remote IP address is monitored. If the remote IP address becomes unavailable, one of the following actions is taken.

- If the action is "wait-recover," packets continue to be sent according to the PBF rule.
- If the action is "fail-over," the firewall uses routing table lookup to determine routing for the duration of this session.

**Table 74. Monitor Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the monitor profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Action | Specify an action to take if the tunnel is not available. If the threshold number of heartbeats is lost, the firewall takes the specified action.<br>• **wait-recover**—Wait for the tunnel to recover; do not take additional action. Packets will continue to be sent according to the PBF rule.<br>• **fail-over**—Traffic will fail over to a backup path, if one is available. The firewall uses routing table lookup to determine routing for the duration of this session.<br>In both cases, the firewall tries to negotiate new IPSec keys to accelerate the recovery. |
| Interval | Specify the time between heartbeats (range 2-10, default 3). |
| Threshold | Specify the number of heartbeats to be lost before the firewall takes the specified action (range 2-100, default 5). |

## Defining Zone Protection Profiles

▶ *Network > Network Profiles > Zone Protection*

Use this page to determine how the firewall responds to attacks from specified security zones. The same profile can be assigned to multiple zones. For an overview of security zones, refer to "Security Zones" on page 151. Flood protection applies to those packets that do not belong to an existing session.

***Note:*** *When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session.*

**Table 75. Zone Protection Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of zone protection profiles when configuring zones. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, and underscores. |
| Description | Enter an optional description for the zone protection profile. |
| **Flood Protection Thresholds - SYN Flood** | |
| Action | Select the action to take in response to a SYN flood attack.<br>• **Random Early Drop**—Causes SYN packets to be dropped to mitigate a flood attack:<br>– When the flow exceeds the **Alert** rate threshold, an alarm is generated.<br>– When the flow exceeds the **Activate** rate threshold, individual SYN packets are dropped randomly to restrict the flow.<br>– When the flow exceeds the **Maximal** rate threshold, all packets are dropped.<br>• **SYN Cookies**—Computes a sequence number for SYN-ACK packets that does not require pending connections to be stored in memory. This is the preferred method. |
| Alert (packets/sec) | Enter the number of SYN packets received by the zone (in a second) that triggers an attack alarm. Alarms can be viewed on the Dashboard (refer to "Using the Dashboard" on page 252) and in the threat log (refer to "Identifying Unknown Applications and Taking Action" on page 276). |
| Activate (packets/sec) | Enter the number of SYN packets received by the zone (in a second) that triggers the action specified. |
| Maximum (packets/sec) | Enter the maximum number of SYN packets able to be received per second. Any number of packets exceeding the maximum will be dropped. |
| **Flood Protection Thresholds - ICMP Flood** | |
| Alert (packets/sec) | Enter the number of ICMP echo requests (pings) received per second that triggers an attack alarm. |
| Activate (packets/sec) | Enter the number of ICMP packets received by the zone (in a second) that causes subsequent ICMP packets to be dropped. |

**Table 75. Zone Protection Profile Settings (Continued)**

| Field | Description |
|---|---|
| Maximum (packets/sec) | Enter the maximum number of ICMP packets able to be received per second. Any number of packets exceeding the maximum will be dropped. |
| **Flood Protection Thresholds - ICMPv6** | |
| Alert (packets/sec) | Enter the number of ICMPv6 echo requests (pings) received per second that triggers an attack alarm. |
| Activate (packets/sec) | Enter the number of ICMPv6 packets received per second for the zone that causes subsequent ICMPv6 packets to be dropped. Metering stops when the number of ICMPv6 packets drops below the threshold |
| Maximum (packets/sec) | Enter the maximum number of ICMPv6 packets able to be received per second. Any number of packets exceeding the maximum will be dropped. |
| **Flood Protection Thresholds - UDP** | |
| Alert (packets/sec) | Enter the number of UDP packets received by the zone (in a second) that triggers an attack alarm. |
| Alert (packets/sec) | Enter the number of UDP packets received by the zone (in a second) that triggers random dropping of UDP packets. The response is disabled when the number of UDP packets drops below the threshold. |
| Maximum (packets/sec) | Enter the maximum number of UDP packets able to be received per second. Any number of packets exceeding the maximum will be dropped. |
| **Flood Protection Thresholds -Other IP** | |
| Alert (packets/sec) | Enter the number of IP packets received by the zone (in a second) that triggers an attack alarm. |
| Activate (packets/sec) | Enter the number of IP packets received by the zone (in a second) that triggers random dropping of IP packets. The response is disabled when the number of IP packets drops below the threshold. Any number of packets exceeding the maximum will be dropped. |
| Maximum (packets/sec) | Enter the maximum number of IP packets able to be received per second. Any number of packets exceeding the maximum will be dropped. |
| **Reconnaissance Protection - TCP Port Scan, UDP Port Scan, Host Sweep** | |
| Interval (sec) | Enter the time interval for port scans and host sweep detection (seconds). |
| Threshold (events) | Enter the number of scanned ports within the specified time interval that will trigger this protection type (events). |
| Action | Enter the action that the system will take in response to this event type:<br>• **Allow**—Permits the port scan of host sweep reconnaissance.<br>• **Alert**—Generates an alert for each scan or sweep that matches the threshold within the specified time interval.<br>• **Block**—Drops all further packets from the source to the destination for the remainder of the specified time interval.<br>• **Block IP**—Drops all further packets for a specified period of time. Choose whether to block source, destination, or source-and-destination traffic and enter a duration (seconds). |

**Table 75. Zone Protection Profile Settings (Continued)**

| Field | Description |
|---|---|
| **IPv6 Drop Packets with** | |
| Type 0 Router Header | Select the check box to drop IPv6 packets that include a Type 0 router header. |
| IPv4 Compatible Address | Select the check box to drop IPv6 packets that include an IPv4-compatible address. |
| Multicast Source Address | Select the check box to drop IPv6 packets that include a multicast source address. |
| Anycast Source Address | Select the check box to drop IPv6 packets that include an anycast source address. |
| **Packet-Based Attack Protection** | |
| **TCP/IP Drop sub tab** | |
| Spoofed IP address | Select the check box to enable protection against IP address spoofing. |
| Fragmented traffic | Discards fragmented IP packets. |
| Mismatched overlapping TCP segment | This setting will cause the firewall to report an overlap mismatch and drop the packet when segment data does not match in these scenarios:<br>• The segment is within another segment.<br>• The segment overlaps with part of another segment.<br>• The segment covers another segment.<br>This protection mechanism uses sequence numbers to determine where packets reside within the TCP data stream. |
| Remove TCP Timestamp | Determines whether the packet has a TCP timestamp in the header and, if it does, strips the timestamp from the header. |
| Reject Non-SYN TCP | Determines whether to reject the packet, if the first packet for the TCP session setup is not a SYN packet:<br>• **global**—Use system-wide setting that is assigned through the CLI.<br>• **yes**—Reject non-SYN TCP.<br>• **no**—Accept non-SYN TCP. Note that allowing non-SYN TCP traffic may prevent file blocking policies from working as expected in cases where the client and/or server connection is not set after the block occurs. |
| Asymmetric Path | Determine whether to drop or bypass packets that contain out of sync ACKs or out of window sequence numbers:<br>• **global**—Use system wide setting that is assigned through the CLI.<br>• **drop**—Drop packets that contain an asymmetric path.<br>• **bypass**—Bypass scanning on packets that contain an asymmetric path. |
| **IP Option Drop** | |
| Strict Source Routing | Discard packets with the Strict Source Routing IP option set. |
| Loose Source Routing | Discard packets with the Loose Source Routing IP option set. |
| Timestamp | Discard packets with the Timestamp IP option set. |
| Record Route | Discard packets with the Record Route IP option set. |
| Security | Discard packets if the security option is defined. |

**Table 75. Zone Protection Profile Settings (Continued)**

| Field | Description |
|---|---|
| Stream ID | Discard packets if the Stream ID option is defined. |
| Unknown | Discard packets if the class and number are unknown. |
| Malformed | Discard packets if they have incorrect combinations of class, number, and length based on RFC 791, 1108, 1393, and 2113. |
| **ICMP Drop sub tab** | |
| ICMP ping ID 0 | Discards packets if the ICMP ping packet has an identifier value of 0. |
| ICMP fragment | Discards packets that consist of ICMP fragments. |
| ICMP large packet (>1024) | Discards ICMP packets that are larger than 1024 bytes. |
| Suppress ICMP TTL expired error | Stop sending ICMP TTL expired messages. |
| Suppress ICMP Frag Needed | Stop sending ICMP fragmentation needed messages in response to packets that exceed the interface MTU and have the do not fragment (DF) bit set. This setting will interfere with the PMTUD process performed by hosts behind the firewall. |
| **IPv6 Drop sub tab** | |
| Type 0 Routing Header | Discard packets with Type 0 Routing header. |
| IPv4 compatible address | Discard packets with IPv4 compatible addresses. |
| Anycast source address | Discard packets with an anycast source address. |
| Needless fragment header | Discard packets with the needless fragment header. |
| MTU in ICMPv6 'Packet Too Big' less than 1280 bytes | Discard packets with MTU in ICMPv6 'Packet Too Big' less than 1280 bytes. |
| Hop-by-Hop extension | Discard packets with the extension header hop-by-hop. |
| Routing extension | Discard packets with the routing extension header. |
| Destination extension | Discard packets with destination extension header. |
| Invalid IPv6 options in extension header | Discard packets with invalid IPv6 option in extension header. |
| Non-zero reserved field | Discard packets if the reserved field has a non-zero value. |
| **ICMPv6 sub tab** | |
| ICMPv6 destination unreachable - require explicit security rule match | Require an explicit security policy match for destination unreachable ICMPv6 errors even when associated with an existing session. |
| ICMPv6 packet too big - require explicit security rule match | Require an explicit security policy match for packet too big ICMPv6 errors even when associated with an existing session. |
| ICMPv6 time exceeded - require explicit security rule match | Require an explicit security policy match for time exceeded ICMPv6 errors even when associated with an existing session. |

**Table 75. Zone Protection Profile Settings (Continued)**

| Field | Description |
|---|---|
| ICMPv6 parameter problem - require explicit security rule match | Require an explicit security policy match for parameter problem ICMPv6 errors even when associated with an existing session. |
| ICMPv6 redirect - require explicit security rule match | Require an explicit security policy match for redirect ICMPv6 messages even when associated with an existing session. |

Chapter 5

# Policies and Security Profiles

This chapter describes how to configure policies and security profiles:

- "Policies" in the next section
- "Security Profiles" on page 211
- "Other Policy Objects" on page 226

## Policies

Policies allow you to control firewall operation by enforcing rules and automatically taking action. The following types of policies are supported:

- Basic security policies to block or allow a network session based on the application, the source and destination zones and addresses, and optionally the service (port and protocol). Zones identify the physical or logical interfaces that send or receive the traffic. Refer to "Security Policies" on page 187.
- Network Address Translation (NAT) policies to translate addresses and ports, as needed. Refer to "NAT Policies" on page 190.
- Policy-based forwarding policies to determine the egress interface used following processing. Refer to "Policy-Based Forwarding Policies" on page 199.
- Decryption policies to specify traffic decryption for security policies. Each policy can specify the categories of URLs for the traffic you want to decrypt. SSH decryption is used to identify and control SSH tunneling in addition to SSH shell access. Refer to "Decryption Policies" on page 202.
- Override policies to override the application definitions provided by the firewall. Refer to "Application Override Policies" on page 205.
- Quality of Service (QoS) policies to determine how traffic is classified for treatment when it passes through an interface with QoS enabled. Refer to "Defining QoS Policies" on page 359.

- Captive portal policies to request authentication of unidentified users. Refer to "Captive Portal Policies" on page 206.
- Denial of service (DoS) policies to protect against DoS attacks and take protective action in response to rule matches. Refer to "DoS Protection Policies" on page 209.

***Note:*** *Shared polices pushed from Panorama are shown in green on the firewall web interface pages and cannot be edited at the device level.*

## Guidelines on Defining Policies

For general guidelines on interacting with the firewall interface, refer to "Using the Firewall Web Interface" on page 23. The following specific guidelines apply when interacting with the pages on the **Policies** tab:

- To apply a filter to the list, select from the **Filter Rules** drop-down list. To add a value to define a filter, click the down-facing arrow for the item and choose **Filter**.
- To view application groups, filters, or container information when creating or viewing Security, PBF, or QoS policies, hold your mouse over the object in the **Application** column, click the down arrow and select **Value**. This allows you to easily view application members directly from the policy without having to navigate to the Object tabs.
- To add a new policy rule, do one of the following:
  - Click **Add** at the bottom of the page.
  - Select a rule on which to base the new rule and click **Clone Rule**, or select a rule by clicking the white space of the rule, and select **Clone Rule** at the bottom of the page (a selected rule has a yellow background). The copied rule, "rule*n*" is inserted below the selected rule, where *n* is the next available integer that makes the rule name unique.
- The order in which rules are listed is the order in which the rules are compared against network traffic. Change the ordering of a rule in either of the following ways:
  - Select the rule and click **Move Up**, **Move Down**, **Move Top**, or **Move Bottom.**
  - Click the down-facing arrow for the rule name and choose **Move**. In the pop-up window, choose a rule and choose whether to move the rule you selected for reordering before or after this rule.

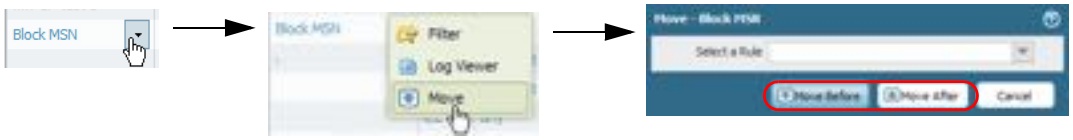


- To enable a rule, select the rule and click **Enable**.

- To show which rules are not currently used, select the **Highlight Unused Rules** check box.

- To display the log for the policy, click the down-facing arrow for the rule name and choose **Log Viewer**.
- For some entries, you can display the current value by clicking the down-facing arrow for the entry and choosing **Value**. You can also edit, filter, or remove certain items directly from the column menu.

- If you have a large number of policies defined, you can use the filter bar to find objects that are used within a policy based on the object name or IP address. The search will also traverse embedded objects to find an address within an address object or address group. In the following screen shot, the IP address 10.8.10.177 was entered in the filter bar and the policy "aaa" is shown. That policy uses an address group object named "aaagroup", which contains the IP address.

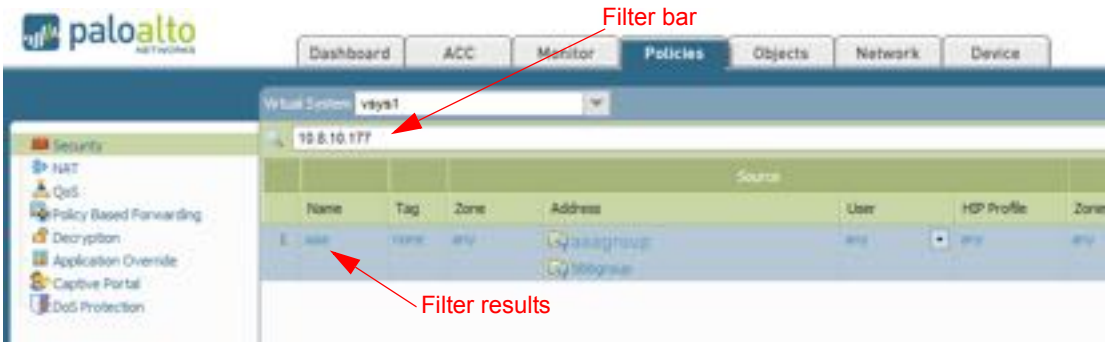

- You can show or hide specific columns from view in any of the **Policies** pages.

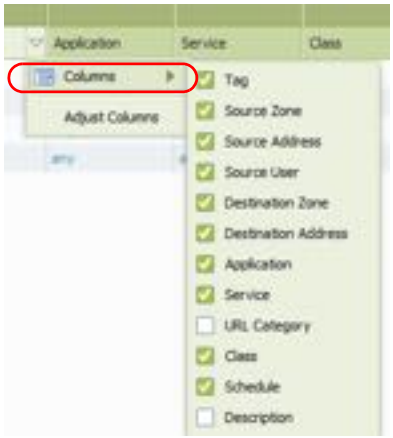


## Specifying Users and Applications for Policies

▶ *Policies > Security*

▶ *Policies > Decryption*

You can restrict security policies to selected users or applications by clicking the **User** or **Application** link on the **Security** or **Decryption** device rules page. For information on restricting rules by application, refer to "Defining Applications" on page 233.

To restrict a policy to selected users, follow these steps:

1. On the **Security** or **Decryption** device rules page, click the underlined link for the source or destination user to open the selection window.

***Note:*** *If you are using a RADIUS server and not the User-ID Agent, the list of users is not displayed, and you must enter user information manually.*

2. Choose the type of rule to apply:
   - **any**—Includes any user in the rule.
   - **pre-logon**—Include remote users that are connected to the network using GlobalProtect, but are not logged into their system. When the Pre-logon option is configured on the Portal for GlobalProtect clients, any user who is not currently logged into their machine will be identified with the username **pre-logon**. You can then create policies for pre-logon users and the machines can be accessed as if they were logged in and connected to the domain.
   - **known-user**—Includes all authenticated users.
   - **unknown**—Includes all unauthenticated users.
   - **Select**—Includes selected users as determined by the selection in this window.
3. To add groups of users, select from the **Available User Groups** check boxes and click **Add User Group**. Alternatively, you can enter text to match one or more groups and click **Add User Group**.

4. To add individual users, enter a search string in the **User** search field and click **Find**. You can then select users and click **Add User**. Alternatively, you can enter individual user names in the **Additional Users** area.
5. Click **OK** to save the selections and update the security or decryption rule.

# Security Policies

Security policies determine whether to block or allow a new network session based on traffic attributes such as the application, source and destination security zones, the source and destination addresses, and the application service (such as HTTP). Security zones are used to group interfaces according to the relative risk of the traffic they carry. For example, an interface connected to the Internet is in an "untrusted" zone, while an interface connected to the internal network is in a "trusted" zone.

***Note:*** *By default, traffic between each pair of security zones is blocked until at least one rule is added to allow traffic between the two zones.*

*Intra-zone traffic is allowed by default and requires an explicit block rule. If a deny all rule is added as the last rule in the policy, intra-zone traffic will be blocked unless otherwise allowed.*

Security policies can be as general or specific as needed. The policy rules are compared against the incoming traffic in sequence, and because the first rule that matches the traffic is applied, the more specific rules must precede the more general ones. For example, a rule for a single application must precede a rule for all applications if all other traffic-related settings are the same.

## Defining Security Policies

▶ *Policies > Security*

Use the **Security** page to define security policy rules. For configuration guidelines, refer to "Guidelines on Defining Policies" on page 184.

**Table 76. Security Policy Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Name | Enter a name to identify the rule (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required. |
| Description | Enter a description for the policy (up to 255 characters). |
| Tag | If you need to tag the policy, click **Add** to specify the tag.<br>A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location. |

**Table 76. Security Policy Settings (Continued)**

| Field | Description |
|---|---|
| **Source Tab** | |
| Source Zone | Click **Add** to choose source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.<br>Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases. |
| Source Address | Click **Add** to add source addresses, address groups, or regions (default is any). Select from the drop-down list, or click the **Address**, **Address Group**, or **Regions** link at the bottom of the drop-down list, and specify the settings. |
| **User Tab** | |
| Source User | Click **Add** to choose the source users or groups of users subject to the policy. |
| HIP Profiles | Click **Add** to choose Host Information Profiles (HIP) to identify users. For information on HIP, refer to "Overview" on page 335. |
| **Destination Tab** | |
| Destination Zone | Click **Add** to choose destination zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.<br>Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases. |
| Destination Address | Click **Add** to add destination addresses, address groups, or regions (default is any). Select from the drop-down list, or click the **Address** link at the bottom of the drop-down list, and specify address settings. |
| **Application Tab** | |
| Application | Select specific applications for the security rule. To define new applications, refer to "Defining Applications" on page 233. To define application groups, refer to "Defining Application Groups" on page 238.<br>If an application has multiple functions, you can select the overall application or individual functions. If you select the overall application, all functions are included, and the application definition is automatically updated as future functions are added.<br>If you are using application groups, filters, or container in the security rule, you can view details on these objects by holding your mouse over the object in the **Application** column, click the down arrow and select **Value**. This allows you to easily view application members directly from the policy without having to navigate to the Object tabs. |

**Table 76. Security Policy Settings (Continued)**

| Field | Description |
|---|---|
| **Service/ URL Category Tab** | |
| Service | Select services to limit to specific TCP and/or UDP port numbers. Choose one of the following from the drop-down list:<br>• **any**—The selected applications are allowed or denied on any protocol or port.<br>• **application-default**—The selected applications are allowed or denied only on their default ports defined by Palo Alto Networks. This option is recommended for allow policies because it prevents applications from running on unusual ports and protocols, which if not intentional, can be a sign of undesired application behavior and usage. Note that when you use this option, the device still checks for all applications on all ports, but with this configuration, applications are only allowed on their default ports/protocols.<br>• **Select**—Click **Add**. Choose an existing service or choose **Service** or **Service Group** to specify a new entry. Refer to "Services" on page 239 and "Service Groups" on page 240. |
| URL Category | Select URL categories for the security rule.<br>• Choose **any** to allow or deny all sessions regardless of the URL category.<br>• To specify a category, click **Add** and select a specific category (including a custom category) from the drop-down list. You can add multiple categories. Refer to "Custom URL Categories" on page 242 for information on defining custom categories. |
| **Actions Tab** | |
| Action Setting | Click **allow** or **deny** to allow or block a new network session for traffic that matches this rule. |
| Profile Setting | To specify the checking done by the default security profiles, select individual Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking, and/or Data Filtering profiles.<br>To specify a profile group, rather than individual profiles, select **Profile Type Group** and then select a profile group from the **Group Profile** drop-down list.<br>To define new profiles or profile groups, click **New** next to the appropriate profile or group (refer to "Security Profile Groups" on page 246). |

**Table 76. Security Policy Settings (Continued)**

| Field | Description |
|---|---|
| Log Setting | Specify any combination of the following options:<br>Log Setting:<br>• To forward the local traffic log and threat log entries to remote destinations, such as Panorama and syslog servers, select a log profile from the **Log Forwarding Profile** drop-down list. Note that the generation of threat log entries is determined by the security profiles. To define new log profiles, click **New** (refer to "Log Forwarding" on page 247).<br>• To generate entries in the local traffic log for traffic that matches this rule, select the following options:<br>– **Log At Session Start**. Generates a traffic log entry for the start of a session (disabled by default).<br>– **Log At Session End**. Generates a traffic log entry for the end of a session (enabled by default).<br>If the session start or end entries are logged, drop and deny entries are also logged. |
| Other Settings | Specify any combination of the following options:<br>• **Schedule**—To limit the days and times when the rule is in effect, select a schedule from the drop-down list. To define new schedules, click **New** (refer to "Schedules" on page 250).<br>• **QoS Marking**—To change the Quality of Service (QoS) setting on packets matching the rule, select IP DSCP or IP Precedence and enter the QoS value in binary or select a predefined value from the drop-down list. For more information on QoS, refer to "Configuring Quality of Service" on page 355.<br>• **Disable Server Response Inspection—**To disable packet inspection from the server to the client, select this check box. This option may be useful under heavy server load conditions. |

# NAT Policies

If you define Layer 3 interfaces on the firewall, you can use Network Address Translation (NAT) policies to specify whether source or destination IP addresses and ports are converted between public and private addresses and ports. For example, private source addresses can be translated to public addresses on traffic sent from an internal (trusted) zone to a public (untrusted) zone.

NAT is also supported on virtual wire interfaces. When performing NAT on virtual wire interfaces, it is recommended that you translate the source address to a different subnet than the one on which the neighboring devices are communicating. Proxy ARP is not supported on virtual wires and so neighboring devices will only be able to resolve ARP requests for IP addresses that reside on the interface of the device on the other end of the virtual wire.

When configuring NAT on the firewall, it is important to note that a security policy must also be configured to allow the NAT traffic. Security policy will be matched based on the post-NAT zone and the pre-NAT IP address.

The firewall supports the following types of address translation:

- **Dynamic IP/Port**—For outbound traffic. Multiple clients can use the same public IP addresses with different source port numbers. Dynamic IP/Port NAT rules allow translation to a single IP address, a range of IP addresses, a subnet, or a combination of

these. In cases where an egress interface has a dynamically assigned IP address, it can be helpful to specify the interface itself as the translated address. By specifying the interface in the dynamic IP/port rule, NAT policy will update automatically to use any address acquired by the interface for subsequent translations.

***Note:*** *Palo Alto Networks Dynamic IP/port NAT supports more NAT sessions than are supported by the number of available IP addresses and ports. The firewall can use IP address and port combinations up to two times (simultaneously) on the PA-200, PA-500, PA-2000 Series and PA-3000 Series, four times on the PA-4020 and PA-5020, and eight times on the PA-4050, PA-4060, PA-5050, and PA-5060 devices when destination IP addresses are unique.*

- **Dynamic IP**—For outbound traffic. Private source addresses translate to the next available address in the specified address range. Dynamic IP NAT policies allow you to specify a single IP address, multiple IPs, multiple IP ranges, or multiple subnets as the translated address pool. If the source address pool is larger than the translated address pool, new IP addresses seeking translation will be blocked while the translated address pool is fully utilized. To avoid this issue, you can specify a fall back pool that will be used if the primary pool runs out of IP addresses.

- **Static IP**—For inbound or outbound traffic. You can use static IP to change the source or the destination IP address while leaving the source or destination port unchanged. When used to map a single public IP address to multiple private servers and services, destination ports can stay the same or be directed to different destination ports.

***Note:*** *You may need to define static routes on the adjacent router and/or the firewall to ensure that traffic sent to a public IP address is routed to the appropriate private address. If the public address is the same as the firewall interface (or on the same subnet), then a static route is not required on the router for that address. When you specify service (TCP or UDP) ports for NAT, the pre-defined HTTP service (service-http) includes two TCP ports: 80 and 8080. To specify a single port, such as TCP 80, you must define a new service.*

The next table summarizes the NAT types. The two dynamic methods map a range of client addresses (M) to a pool (N) of NAT addresses, where M and N are different numbers. N can also be 1. Dynamic IP/Port NAT differs from Dynamic IP NAT in that the TCP and UDP source ports are not preserved in Dynamic IP/Port, whereas they are unchanged with Dynamic IP NAT. There are also differing limits to the size of the translated IP pool, as noted below.

With Static IP NAT, there is a one-to-one mapping between each original address and its translated address. This can be expressed as 1-to-1 for a single mapped IP address, or M-to-M for a pool of many one-to-one, mapped IP addresses.

**Table 77. NAT Types**

| PAN-OS NAT Type | Source Port Stays the Same | Destination Port Can Change | Mapping Type | Size of Translated Address Pool |
|---|---|---|---|---|
| Dynamic IP/ Port | No | No | Many-to-1<br>M-to-N | Up to 254 consecutive addresses |
| Dynamic IP | Yes | No | M-to-N | Up to 32k consecutive addresses |
| Static IP | Yes | No | 1-to-1<br>M-to-M<br>MIP | Unlimited |
| | Optional | | 1-to-Many VIP PAT | |

## Determining Zone Configuration in NAT and Security Policy

NAT rules must be configured to use the zones associated with pre-NAT IP addresses configured in the policy. For example, if you are translating traffic that is incoming to an internal server (which is reached via a public IP by Internet users), it is necessary to configure the NAT policy using the zone in which the public IP address resides. In this case, the source and destination zones would be the same. As another example, when translating outgoing host traffic to a public IP address, it is necessary to configure NAT policy with a source zone corresponding to the private IP addresses of those hosts. The pre-NAT zone is required because this match occurs before the packet has been modified by NAT.

Security policy differs from NAT policy in that post-NAT zones must be used to control traffic. NAT may influence the source or destination IP addresses and can potentially modify the outgoing interface and zone. When creating security policies with specific IP addresses, it is important to note that pre-NAT IP addresses will be used in the policy match. Traffic subject to NAT must be explicitly permitted by the security policy when that traffic traverses multiple zones.

## NAT Rule Options

The firewall supports no-NAT rules and bi-directional NAT rules.

### No-NAT Rules

No-NAT rules are configured to allow exclusion of IP addresses defined within the range of NAT rules defined later in the NAT policy. To define a no-NAT policy, specify all of the match criteria and select **No Source Translation** in the source translation column.

### Bi-directional NAT Rules

The bi-directional setting in static source NAT rules implicitly creates a destination NAT rule for traffic to the same resources in the reverse direction. In this example, two NAT rules are used to create a source translation for outgoing traffic from IP 10.0.1.10 to public IP 3.3.3.1 and a destination translation for traffic destined for public IP 3.3.3.1 to private IP 10.0.1.10. This pair of rules can be simplified by configuring only the third NAT rule using the bi-directional feature.

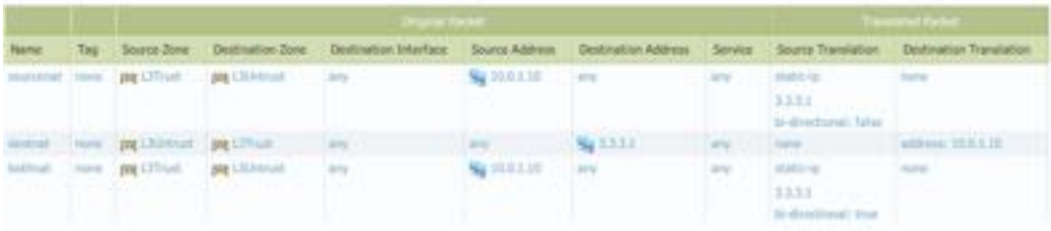


**Figure 17. Bi-Directional NAT Rules**

## Defining Network Address Translation Policies

▶ *Policies > NAT*

NAT address translation rules are based on the source and destination zones, the source and destination addresses, and the application service (such as HTTP). Like security policies, the NAT policy rules are compared against the incoming traffic in sequence, and the first rule that matches the traffic is applied.

As needed, add static routes to the local router so that traffic to all public addresses is routed to the firewall. You may also need to add static routes to the receiving interface on the firewall to route traffic back to the private address (refer to "Firewall Interfaces" on page 127). For configuration guidelines, refer to "Guidelines on Defining Policies" on page 184.

**Table 78. NAT Rule Settings**

| Field | Description |
|---|---|
| Name | Change the default rule name and/or enter a rule description. |
| Description | Enter a description for the policy (up to 255 characters). |
| NAT Type | Specify **ipv4** for NAT between IPv4 addresses, or **nat64** translation between IPv6 and IPv4 addresses.<br>***Note:*** *You cannot combine IPv4 and IPv6 address ranges in a single NAT rule.* |
| Tag | If you need to tag the policy, click **Add** to specify the tag.<br>A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location. |
| **Original Packet** | |
| Source Zone<br>Destination Zone | Select one or more source and destination zones for the original (non-NAT) packet (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.<br>Multiple zones can be used to simplify management. For example, you can configure settings so that multiple internal NAT addresses are directed to the same external IP address. |
| Destination Interface | Specify the type of interface for translation. The destination interface can be used to translate IP addresses differently in the case where the network is connected to two ISPs with different IP address pools. |
| Source Address<br>Destination Address | Specify a combination of source and destination addresses for which the source or destination address must be translated. |
| Service | Specify the services for which the source or destination address is translated. To define new service groups, refer to "Service Groups" on page 240. |

**Table 78. NAT Rule Settings (Continued)**

| Field | Description |
|---|---|
| **Translated Packet** | |
| Source Translation | Enter an IP address or address range (address1-address2) that the source address is translated to, and select a dynamic or static address pool. The size of the address range is limited by the type of address pool:<br>• **Dynamic IP And Port**—Address selection is based on a hash of the source IP address. For a given source IP address, the firewall will use the same translated source address for all sessions. Dynamic IP and Port source NAT supports approximately 64k concurrent sessions on each IP address in the NAT pool. On some platforms, over-subscription is supported, which will allow a single IP to host more than 64k concurrent sessions. See the note on page 191 for more details.<br>• **Dynamic IP**—The next available address in the specified range is used, but the port number is unchanged. Up to 32k consecutive IP addresses are supported. A dynamic IP pool can contain multiple subnets, so you can translate your internal network addresses to two or more separate public subnets.<br>– **Advanced (Fall back Dynamic IP Translation)**—Use this option to create a fall back pool that will perform IP and port translation and will be used if the primary pool runs out of addresses. You can define addresses for the pool by using the Translated Address option or the Interface Address option, which is for interfaces that receive an IP address dynamically. When creating a fall back pool, make sure addresses do not overlap with addresses in the primary pool.<br>• **Static IP**—The same address is always used, and the port is unchanged. For example, if the source range is 192.168.0.1-192.168.0.10 and the translation range is 10.0.0.1-10.0.0.10, address 192.168.0.2 is always translated to 10.0.0.2. The address range is virtually unlimited.<br>• **None**—Translation is not performed. |
| Destination Translation | Enter an IP address or range of IP addresses and a translated port number (1 to 65535) that the destination address and port number are translated to. If the **Translated Port** field is blank, the destination port is not changed. Destination translation is typically used to allow an internal server, such as an email server, to be accessed from the public network. |

## NAT Policy Examples

The following NAT policy rule translates a range of private source addresses (10.0.0.1 to 10.0.0.100 in the "L3Trust" zone) to a single public IP address (200.10.2.100 in the "L3Untrust" zone) and a unique source port number (dynamic source translation). The rule applies only to traffic received on a Layer 3 interface in the "L3Trust" zone that is destined for an interface in the "L3Untrust" zone. Because the private addresses are hidden, network sessions cannot be initiated from the public network. If the public address is not a firewall interface address (or on the same subnet), the local router requires a static route to direct return traffic to the firewall.

Security policy must be explicitly configured to permit traffic matching this NAT rule. Create a security policy with source/destination zones and source/destination addresses matching the NAT rule.

**Figure 18. Dynamic Source Address Translation**

In the following example, the first NAT rule translates the private address of an internal mail server to a static public IP address. The rule applies only to outgoing email sent from the "L3Trust" zone to the "L3Untrust" zone. For traffic in the reverse direction (incoming email), the second rule translates the destination address from the server's public address to its private address. Rule2 uses "L3Untrust" for the source and destination zones because NAT policy is based on the pre-NAT address zone. In this case, that pre-NAT address is a public IP address and is therefore in the "L3Untrust" zone.

**Figure 19. Static Source and Destination Address Translation**

In both examples, if the public address is not the address of the firewall's interface (or on the same subnet), you must add a static route to the local router to route traffic to the firewall.

## NAT64

NAT64 is used to translate source and destination IP headers between IPv6 and IPv4 addresses. It allows IPv6 clients to access IPv4 servers and allows IPv4 clients to access IPv6 servers. There are three main transition mechanisms defined by the IETF: dual-stack, tunneling, and translation. When you have IPv4-only and IPv6-only networks and communication is required, you must use translation.

When using NAT64 policies on Palo Alto Networks firewall, it's required that you have a third-party DNS64 solution in place to decouple the DNS query function from the NAT function.

The following NAT64 features are supported:

- Stateful NAT64, which allows preserving of IPv4 addresses so that one IPv4 address can map to multiple IPv6 addresses. An IPv4 address can also be shared with NAT44. In contrast, stateless NAT64 maps one IPv4 address to one IPv6 address.
- IPv4-initiated communication translation. The static binding for IPv4 maps an IPv4 address/port number to an IPv6 IP address. PAN-OS also supports port rewrite that allows you to preserve even more IPv4 addresses.
- Allows translation for /32, /40, /48, /56, /64, and /96 subnets.
- Support for multiple prefixes. You can assign one NAT64 prefix per rule.
- Does not require you to reserve a pool of IPv4 address specifically for NAT64. Therefore, it is possible to use a single IP address to do NAT44 and NAT64.
- Supports hairpinning (NAT U-Turn) and it can prevent hairpinning loop attacks.

- Supports the translation of TCP/UDP/ICMP packets as per RFC, but also other protocols without ALG (best effort). For example, a GRE packet could be translated. This translation has the same limitation as NAT44.
- Supports PMTUD (path MTU discovery) and it updates the MSS (Maximum Segment Size) for TCP.
- Allows configuration of the IPv6 MTU setting. The default value is 1280, which is the minimum MTU for IPv6 traffic. This setting is configured in **Device > Setup > Sessions** tab under **Session Settings**.
- Translates length attribute between IPv4 and IPv6.
- Supported on Layer 3 interfaces and subinterfaces, tunnel, and VLAN interfaces.

## NAT64 Examples

You can configure two types of translation with the firewall: IPv6-initiated communication, which is similar to source NAT in IPv4, and IPv4-initiated communication to an IPv6 server, which is similar to destination NAT in IPv4.

### IPv6-initiated Communication

In this type of translation, the destination IPv6 address in the NAT rule is a prefix following the RFC 6052 format (/32, /40,/48,/56,/64, and /96). The destination IPv6 address netmask in the rule would be used to extract the IPv4 address. The source translation needs to have "Dynamic IP and Port" in order to implement a stateful NAT64. The IPv4 address set as the source is configured the same way as a NAT44 source translation. The destination translation field is not set. However, a destination translation must be done since the address is extracted from the IPv6 address in the packet. It uses prefix defined in the destination IP matching criteria. You should note that in a /96 prefix, it is the last 4 octets, but the location of the IPv4 address would be different if the prefix is not /96.

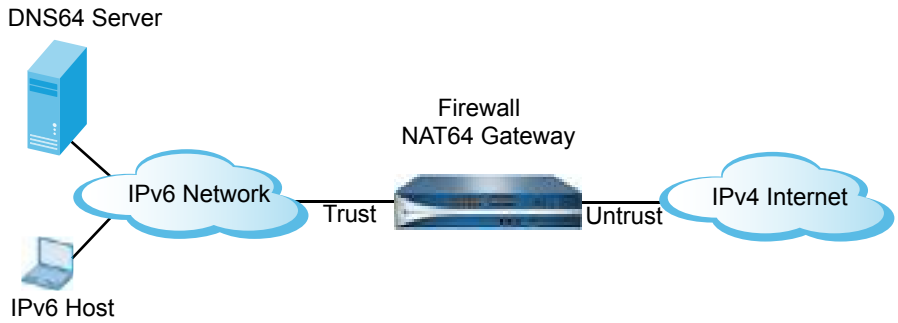


**Figure 20. NAT64 IPv6 Client to IPv4 Network Diagram**

The following table describes the values needed in this NAT64 policy.

**Table 79.**

| Source IP | Destination IP | Source Translation | Destination Translation |
|---|---|---|---|
| Any/IPv6 address | NAT64 IPv6 prefix with RFC6052 compliant netmask | Dynamic IP and port mode (Use IPv4 address) | None<br>(Extracted from the destination IPv6 address) |

### IPv4-initiated Communication

The IPv4 address is the address that maps to the IPv6 address and you use static IP mode in the source translation. The source would be set in an IPv6 prefix as defined in RFC6052 and is appended to the IPv4 source address. The destination address is the IP address set in the destination translation column. It is possible to rewrite the destination port. This method allows a single IP address to share multiple IPv6 servers through the port through a static mapping.

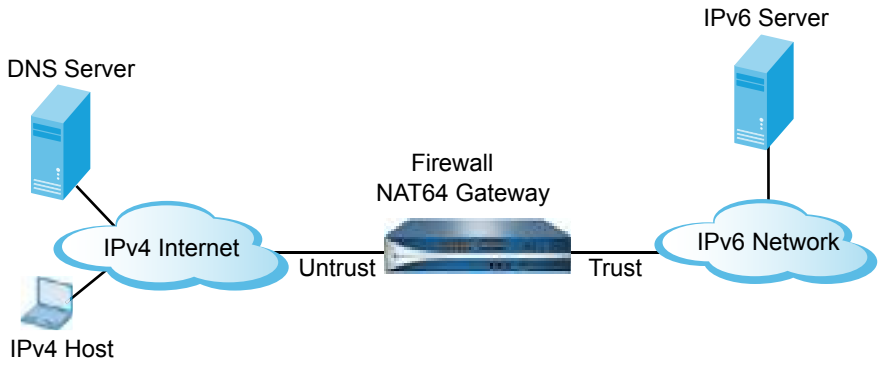


**Figure 21. NAT64 IPv4 Internet to IPv6 Customer Network Diagram**

The following table describes the values needed in this NAT64 policy.

**Table 80. IPv4-initiated Values**

| Source IP | Destination IP | Source Translation | Destination Translation |
|---|---|---|---|
| Any/IPv4 address | IPv4 address | Static IP mode<br>(IPv6 prefix in RFC 6052 format) | Single IPv6 address (actual server IP address)<br>Note: You could specify a server port re-write. |

The packet processing engine of the firewall must do a route lookup to find the destination zone prior to looking at the NAT rule. In NAT64, it is important to address the reachability of the NAT64 prefix for the destination zone assignment since the NAT64 prefix should not be routable by the NAT64 gateway. It is very likely that the NAT64 prefix would hit the default route, or be dropped because there is no route. You can setup a tunnel interface with no termination point since this type of interface will act like a loopback port and accept other netmasks besides /128. You apply the NAT64 prefix to the tunnel and apply the appropriate

zone in order to ensure that IPv6 traffic with NAT64 prefix is assigned to the proper destination zone. It would also have the advantage to drop the IPv6 traffic with NAT64 prefix if the NAT64 rule is not matched.

### IETF Scenarios for IPv4/IPv6 Translation

There are six NAT64 based scenarios defined by the IETF in RFC 6144. The Palo Alto Networks firewall will support all but one of these scenarios, as described in the following table.

**Table 81. Summary of IETF Scenario Implementations for Using PAN-OS**

| Scenario | Source IP | Destination IP | Source Translation | Destination Translation |
|---|---|---|---|---|
| IPv6 Network to the IPv4 Internet | Any/IPv6 address | NAT64 IPv6 prefix with RFC 6052 compliant netmask. | Dynamic IP and port mode.<br>Use *Public* IPv4 address | None<br>(extracted from destination IPv6 address) |
| The IPv4 Internet to an IPv6 Network | Any/IPv4 address | Single IPv4 address | Static IP mode.<br>IPv6 Prefix in RFC 6052 format | Single IPv6 address |
| The IPv6 Internet to an IPv4 Network | Any/IPv6 address | IPV6 globally routable prefix with RFC 6052 compliant netmask. | Dynamic IP and port.<br>Use *Private* IPv4 address | None<br>(extracted from destination IPv6 address) |
| IPv4 network to IPv6 Internet | Not currently supported | | | |
| IPv4 network to IPv6 network | Any/IPv4 address | Single IPv4 address | Static IP mode.<br>IPv6 Prefix in RFC 6052 format | Single IPv6 address |
| IPv6 network to IPv4 network | Any/IPv6 address | NAT64 IPV6 prefix with RFC 6052 compliant netmask. | Dynamic IP and port.<br>Use Private IPv4 address | None<br>(extracted from destination IPv6 address) |

# Policy-Based Forwarding Policies

▶ ***Policies > Policy Based Forwarding***

Normally, when traffic enters the firewall, the ingress interface virtual router dictates the route that determines the outgoing interface and destination security zone based on destination IP address. With policy-based forwarding (PBF), you can specify other information to determine the outgoing interface, including source zone, source address, source user, destination address, destination application, and destination service. The initial session on a given destination IP address and port that is associated with an application will not match an application-specific rule and will be forwarded according to subsequent PBF rules (that do not specify an application) or the virtual router's forwarding table. All subsequent sessions on

that destination IP address and port for the same application will match an application-specific rule. To ensure forwarding through PBF rules, application-specific rules are not recommended.

When necessary, PBF rules can be used to force traffic through an additional virtual system using the Forward-to-VSYS forwarding action. In this case, it is necessary to define an additional PBF rule that will forward the packet from the destination virtual system out through a particular egress interface on the firewall.

For configuration guidelines, refer to "Guidelines on Defining Policies" on page 184.

**Table 82. Policy-Based Forwarding Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Name | Enter a name to identify the rule (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required. |
| Description | Enter a description for the policy (up to 255 characters). |
| Tag | If you need to tag the policy, click **Add** to specify the tag.<br>A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location. |
| **Source Tab** | |
| Source Zone | To choose source zones (default is any), click **Add** and select from the drop-down list. To define new zones, refer to "Defining Security Zones" on page 151.<br>Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases.<br>**Note:** Only Layer 3 type zones are supported for policy-based forwarding. |
| Source Address | Click **Add** to add source addresses, address groups, or regions (default is any). Select from the drop-down list, or click the **Address**, **Address Group**, or **Regions** link at the bottom of the drop-down list, and specify the settings. |
| Source User | Click **Add** to choose the source users or groups of users subject to the policy. |
| **Destination/ Application/Service Tab** | |
| Destination Address | Click **Add** to add destination addresses, address groups, or regions (default is any). Select from the drop-down list, or click the **Address**, **Address Group**, or **Regions** link at the bottom of the drop-down list, and specify the settings. |

**Table 82. Policy-Based Forwarding Settings (Continued)**

| Field | Description |
|---|---|
| Application | Select specific applications for the PBF rule. To define new applications, refer to "Defining Applications" on page 233. To define application groups, refer to "Defining Application Groups" on page 238.<br>If you are using application groups, filters, or container in the PBF rule, you can view details on these objects by holding your mouse over the object in the **Application** column, click the down arrow and select **Value**. This enables you to easily view application members directly from the policy without having to go to the Object tabs. |
| Service | Specify the services for which the source or destination address is translated. To define new service groups, refer to "Service Groups" on page 240. |
| **Forwarding Tab** | |
| Action | Select one of the following options:<br>• **Forward**—Specify the next hop IP address and egress interface (the interface that the packet takes to get to the specified next hop).<br>• **Forward To VSYS**—Choose the virtual system to forward to from the drop-down list.<br>• **Discard**—Drop the packet.<br>• **No PBF**—Do not alter the path that the packet will take. |
| Egress Interface | Specify the firewall interface for forwarding traffic from the firewall. |
| Next Hop | Specify the IP address of the next forwarding stop. |
| Monitor | To monitor the forwarding actions, select **Monitor** and specify the following settings:<br>• **Profile**—Choose a profile from the drop-down list.<br>• **Disable if unreachable**—Select this check box if you want to ignore this rule for all new sessions when the next hop router is unreachable.<br>• **IP Address**—Specify the IP address to which ping messages are sent periodically to determine the state of the policy based forwarding rule. |

**Table 82. Policy-Based Forwarding Settings (Continued)**

| Field | Description |
|---|---|
| Enforce Symmetric Return | Select this option to allow virtual routers on the firewall to enforce symmetric return for the PBF rule. Click **Add** if you want to enter the next-hop address(es) that will be used for the forwarding.<br>This option will circumvent the route lookup process for return traffic; the firewall will use the original incoming interface as the egress interface. If the source IP is in the same subnet as the incoming interface on the firewall, symmetric return will not take effect.<br>This feature is useful when you have servers accessible through two ISP connections (on different ingress interfaces) and the return traffic must be routed through the ISP that original routed the sessions.<br>The following should be noted when using this feature:<br>• If an interface has multiple PBF rules, only one rule can enforce symmetric return.<br>• You can use this option for all L3 interfaces, except loopback interfaces. You can also use interfaces that have the IP address assigned dynamically (DHCP and PPoE).<br>• The source must be an interface, not a zone.<br>• Next-hop address list is not supported for tunnel and PPoE interfaces.<br>• Up to 8 next-hop addresses can be defined per PBF rule. |
| Schedule | To limit the days and times when the rule is in effect, select a schedule from the drop-down list. To define new schedules, refer to "Schedules" on page 250. |

# Decryption Policies

▶ *Policies > Decryption*

You can configure the firewall to decrypt traffic for visibility, control, and granular security. Decryption policies can apply to Secure Sockets Layer (SSL) and Secure Shell (SSH) traffic. With the SSH option, you can selectively decrypt outbound and inbound SSH traffic to assure that secure protocols are not being used to tunnel disallowed applications and content. You can also apply decryption profiles to your policies to block and control various aspects of SSL traffic. For more information, refer to "Decryption Profiles" on page 248.

Each decryption policy specifies the categories of URLs to decrypt or not decrypt. App-ID and the Antivirus, Vulnerability, Anti-spyware, URL Filtering, and File-blocking profiles are applied to decrypted traffic before it is re-encrypted as traffic exits the device. End-to-end security between clients and servers is maintained, and the firewall acts as a trusted third party during the connection. No decrypted traffic leaves the device.

The firewall inspects traffic, regardless of the protocols that are encapsulated. Decryption policies can be as general or specific as needed. The policy rules are compared against the traffic in sequence, so more specific rules must precede the more general ones.

***Note:*** *Refer to the Palo Alto Networks tech note, Controlling SSL Decryption, for instructions on managing SSL certificates to avoid certificate mismatch errors, and SSL Decryption Certificates for guidelines on how to develop policies to handle non-standard SSL implementations.*

SSL forward proxy decryption requires the configuration of a trusted certificate that will be presented to the user if the server to which the user is connecting possesses a certificate signed by a CA trusted by the firewall. To configure this certificate, create a certificate on the **Device > Certificate Management > Certificates** page and then click the name of the certificate and check the **Forward Trust Certificate** check box. Refer to "Importing, Exporting and Generating Security Certificates" on page 86.

***Note:*** *Certain applications will not function if they are decrypted by the firewall. To prevent this from occurring, PAN-OS will not decrypt the SSL traffic for these applications and the decryption rule settings will not apply. For a list of these applications, refer to support article located at https://live.paloaltonetworks.com/docs/DOC-1423.*

.

**Table 83. Decryption Rule Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Name | Enter a name to identify the rule (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required. |
| Description | Enter a description for the rule (up to 255 characters). |
| Tag | If you need to tag the policy, click **Add** to specify the tag.<br>A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location. |
| **Source Tab** | |
| Source Zone | Click **Add** to choose source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.<br>Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases. |
| Source Address | Click **Add** to add source addresses, address groups, or regions (default is any). Select from the drop-down list, or click the **Address**, **Address Group**, or **Regions** link at the bottom of the drop-down list, and specify the settings. Select the **Negate** check box to choose any address except the configured ones. |

**Table 83. Decryption Rule Settings (Continued)**

| Field | Description |
|---|---|
| Source User | Click **Add** to choose the source users or groups of users subject to the policy. If you choose the **pre-logon** option, this will include remote users that are connected to the network using GlobalProtect, but are not logged into their system. When the Pre-logon option is configured on the Portal for GlobalProtect clients, any user who is not currently logged into their machine will be identified with the username **pre-logon**. You can then create policies for pre-logon users and the machines can be accessed as if they were logged in and connected to the domain. |
| **Destination Tab** | |
| Destination Zone | Click **Add** to choose destination zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.<br>Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases. |
| Destination Address | Click **Add** to add destination addresses, address groups, or regions (default is any). Select from the drop-down list, or click the **Address**, **Address Group**, or **Regions** link at the bottom of the drop-down list, and specify the settings. Select the **Negate** check box to choose any address except the configured ones. |
| **URL Category Tab** | Select URL categories for the decryption rule.<br>• Choose **any** to match any sessions regardless of the URL category.<br>• To specify a category, click **Add** and select a specific category (including a custom category) from the drop-down list. You can add multiple categories. Refer to "Custom URL Categories" on page 242 for information on defining custom categories. |
| **Options Tab** | |
| Action | Select **decrypt** or **no-decrypt** for the traffic. |
| Type | Select the type of traffic to decrypt from the drop-down list:<br>• **SSL Forward Proxy**—Specifies that the policy will decrypt client traffic destined for an external server.<br>• **SSH Proxy**—Specifies that the policy will decrypt SSH traffic. This option allows you to control SSH tunneling in policies by specifying the ssh-tunnel App-ID.<br>• **SSL Inbound Inspection**—Specifies that the policy will decrypt SSL inbound inspection traffic. |
| Decryption Profile | Select an existing decryption profile, or create a new decryption profile. Refer to "Decryption Profiles" on page 248. |

# Application Override Policies

To change how the firewall classifies network traffic into applications, you can specify application override policies. For example, if you want to control one of your custom applications, an application override policy can be used to identify traffic for that application according to zone, source and destination address, port, and protocol. If you have network applications that are classified as "unknown," you can create new application definitions for them (refer to "Defining Applications" on page 233).

Like security policies, application override policies can be as general or specific as needed. The policy rules are compared against the traffic in sequence, so the more specific rules must precede the more general ones.

## Custom Application Definition with Application Override

Because the App-ID engine in PAN-OS classifies traffic by identifying the application-specific content in network traffic, the custom application definition cannot simply use a port number to identify an application. The application definition must also include traffic (restricted by source zone, source IP address, destination zone, and destination IP address).

To create a custom application with application override:

1. Define the custom application. Refer to "Defining Applications" on page 233. It is not required to specify signatures for the application if the application is used only for application override rules.
2. Define an application override policy that specifies when the custom application should be invoked. A policy typically includes the IP address of the server running the custom application and a restricted set of source IP addresses or a source zone.

## Defining Application Override Policies

▶ ***Policies*** **>** ***Application Override***

After creating a new rule, configure the rule by clicking the current field values and specifying the appropriate information, as described in the following table.

**Table 84. Application Override Rule Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Name | Enter a name to identify the rule (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required. |
| Description | Enter a description for the rule (up to 255 characters). |
| Tag | If you need to tag the policy, click **Add** to specify the tag.<br>A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location. |

**Table 84. Application Override Rule Settings (Continued)**

| Field | Description |
|---|---|
| **Source Tab** | |
| Source Zone | Click **Add** to choose source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.<br>Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases. |
| Source Address | Click **Add** to add source addresses, address groups, or regions (default is any). Select from the drop-down list, or click the **Address**, **Address Group**, or **Regions** link at the bottom of the drop-down list, and specify the settings. Select the **Negate** check box to choose any address except the configured ones. |
| **Destination Tab** | |
| Destination Zone | Click **Add** to choose destination zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.<br>Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases. |
| Destination Address | Click **Add** to add destination addresses, address groups, or regions (default is any). Select from the drop-down list, or click the **Address**, **Address Group**, or **Regions** link at the bottom of the drop-down list, and specify the settings. Select the **Negate** check box to choose any address except the configured ones. |
| **Protocol/Application Tab** | |
| Protocol | Select the protocol for which the application can be overridden. |
| Port | Enter the port number (0 to 65535) or range of port numbers (port1-port2) for the specified destination addresses. Multiple ports or ranges must be separated by commas. |
| Application | Select the override application for traffic flows that match the above rule criteria. When overriding to a custom application, there is no threat inspection that is performed. The exception to this is when you override to a pre-defined application that supports threat inspection.<br>To define new applications, refer to "Defining Applications" on page 233). |

# Captive Portal Policies

You can set up and customize a captive portal to direct user authentication by way of an authentication profile, an authentication sequence, or a certificate profile. Captive portal is used in conjunction with the User-ID Agent to extend user identification functions beyond the Active Directory domain. Users are directed to the portal and authenticated, thereby creating a user-to-IP address mapping.

## Defining Captive Portal Policies

▶ *Policies > Captive Portal*

Before you define captive portal policies, enable captive portal and configure captive portal settings on the **User Identification** page, as described in "Configuring the Firewall for User Identification" on page 286.

After doing so, configure capture portal policies by specifying the following information.

**Table 85. Captive Portal Rule Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the rule (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required. |
| Description | Enter a description for the rule (up to 255 characters). |
| Tag | If you need to tag the policy, click **Add** to specify the tag.<br>A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location. |
| **Source Tab** | |
| Source | Specify the following information:<br>• Choose a source zone if the policy needs to be applied to traffic coming from all interfaces in a given zone. Click **Add** to specify multiple interfaces or zones.<br>• Specify the **Source Address** setting to apply the captive portal policy for traffic coming from specific source addresses. Select the **Negate** check box to choose any address except the configured ones. Click **Add** to specify multiple interfaces or zones. |
| **Destination Tab** | |
| Destination | Specify the following information:<br>• Choose a destination zone if the policy needs to be applied to traffic to all interfaces in a given zone. Click **Add** to specify multiple interfaces or zones.<br>• Specify the **Destination Address** setting to apply the captive portal policy for traffic to specific destination addresses. Select the **Negate** check box to choose any address except the configured ones. Click **Add** to specify multiple interfaces or zones. |

**Table 85. Captive Portal Rule Settings (Continued)**

| Field | Description |
|---|---|
| **Service/ URL Category Tab** | |
| Service | Select services to limit to specific TCP and/or UDP port numbers. Choose one of the following from the drop-down list:<br>• **any**—The selected services are allowed or denied on any protocol or port.<br>• **default**—The selected services are allowed or denied only on the default ports defined by Palo Alto Networks. This option is recommended for allow policies.<br>• **Select**—Click **Add**. Choose an existing service or choose **Service** or **Service Group** to specify a new entry. Refer to "Services" on page 239 and "Service Groups" on page 240. |
| URL Category | Select URL categories for the captive portal rule.<br>• Choose **any** to apply the actions specified on the **Service/Action** tab regardless of the URL category.<br>• To specify a category, click **Add** and select a specific category (including a custom category) from the drop-down list. You can add multiple categories. Refer to "Custom URL Categories" on page 242 for information on defining custom categories. |
| **Service/Action Tab** | |
| Action Setting | Choose an action to take:<br>• **web-form**—Present a captive portal page for the user to explicitly enter authentication credentials.<br>• **no-captive-portal**—Allow traffic to pass without presenting a captive portal page for authentication.<br>• **browser-challenge**—Open an NT LAN Manager (NTLM) authentication request to the user's web browser. The web browser will respond using the user's current login credentials. |

# DoS Protection Policies

DoS protection policies allow you to control the number of sessions between interfaces, zones, addresses, and countries based on aggregate sessions or source and/or destination IP addresses.

## Defining DoS Protection Policies

▶ *Policies > DoS Protection*

Use this page to add, edit, or delete DoS protection policy rules. To add a policy rule, click **Add** and then complete the following fields:

**Table 86. DoS Rule Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Name | Enter a name to identify the rule (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required. |
| Description | Enter a description for the rule (up to 255 characters). |
| Tag | If you need to tag the policy, click **Add** to specify the tag.<br>A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location. |
| **Source Tab** | |
| Source | Specify the following information:<br>• Choose **Interface** from the **Type** drop-down list to apply the DoS policy to traffic coming from an interface or a group of interfaces. Choose **Zone** if the DoS policy needs to be applied to traffic coming from all interfaces in a given zone. Click **Add** to specify multiple interfaces or zones.<br>• Specify the **Source Address** setting to apply the DoS policy for traffic coming from specific source addresses. Select the **Negate** check box to choose any address except the configured ones. Click **Add** to specify multiple addresses.<br>• Specify the **Source User** setting to apply the DoS policy for traffic from specific users. Click **Add** to specify multiple users. |

**Table 86. DoS Rule Settings (Continued)**

| Field | Description |
|---|---|
| **Destination Tab** | |
| Destination | Specify the following information:<br>• Choose **Interface** from the **Type** drop-down list to apply the DoS policy to traffic coming from an interface or a group of interfaces. Choose **Zone** if the DoS policy needs to be applied to traffic coming from all interfaces in a given zone. Click **Add** to specify multiple interfaces or zones.<br>• Specify the **Destination Address** setting to apply the DoS policy for traffic to specific destination addresses. Select the **Negate** check box to choose any address except the configured ones. Click **Add** to specify multiple addresses. |
| **Option/Protection Tab** | |
| Service | Select from the drop-down list to apply the DoS policy to only the configured services. |
| Action | Choose the action from the drop-down list:<br>• **Deny**—Drop all traffic.<br>• **Allow**—Permit all traffic.<br>• **Protect**—Enforce protections supplied in the thresholds that are configured as part of the DoS profile applied to this rule. |
| Schedule | Select a pre-configured schedule from the drop-down list to apply the DoS rule to a specific date/time. |
| Log Forwarding | If you want to trigger forwarding of threat log entries to an external service—such as a syslog server or Panorama—select a log forwarding profile from the drop-down or click **Profile** to create a new one. Note that only traffic that matches an action in the rule will be logged and forwarded. |
| Aggregate | Select a DoS protection profile from the drop-down list to determine the rate at which you want to take action in response to DoS threats. The aggregate setting applies to the total of all traffic from the specified source to specified destination. |
| Classified | Select the check box and specify the following:<br>• **Profile**—Select the profile from the drop-down list.<br>• **Address**—Select whether to apply the rule to the source, destination, or source and destination IP addresses.<br>If a classified profile is specified, the profile limitations are applied to a source IP address, destination IP address, or source and destination IP address pair. For example, you could specify a classified profile with a session limit of 100 and specify an **Address** setting of "source" in the rule. The result would be a limit of 100 sessions at any given time for that particular source IP address. |

# Security Profiles

Each security policy can include specification of one or more security profiles, which provide additional protection and control.

You can also add threat exceptions to Anti-spyware and Vulnerability profiles. To make management of threat exceptions easier, you can add threat exceptions directly from the Monitor > Logs > Threat list. To add a threat exception to a profile, locate the threat in the Threat log and then click on the threat name. The Threat Details dialogue will appear, select one or more profiles in the left pane and click OK. This will add a threat exception for that signature in the selected threat profiles with the action of allow. To add a threat exception for specific IP addresses, add the IP addresses in the right pane and click OK. This will add a threat exception with the IP addresses added as a filter on the threat exception.

The following profile types are available:

- Antivirus profiles to protect against worms and viruses or block spyware downloads. Refer to "Antivirus Profiles" in the next section.
- Anti-spyware profiles to block attempts by spyware to access the network. Refer to "Anti-spyware Profiles" on page 213.
- Vulnerability protection profiles to stop attempts to exploit system flaws or gain unauthorized access to systems. Refer to "Vulnerability Protection Profiles" on page 215.
- URL filtering profiles to restrict access to specific web sites and web site categories. Refer to "URL Filtering Profiles" on page 217.
- File blocking profiles to block selected file types. Refer to "File Blocking Profiles" on page 220.
- Data filtering profiles that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall. Refer to "Data Filtering Profiles" on page 223.

In additional to individual profiles, you can create profile groups to combine profiles that are often applied together.

***Note:*** *You cannot delete a profile that is used in a security policy.*

You can choose from the following actions when defining antivirus and anti-spyware profiles.

- **Default**—Takes the default action that is specified internally in the signature for each threat.
- **Alert**—Generates an alert for each application traffic flow. The alert is saved in the threat log.
- **Block**—Drops the application traffic.
- **Allow**—Permits the application traffic.

The following actions are available when defining custom Spyware and Vulnerability objects:

- **Alert**—Generates an alert for each application traffic flow. The alert is saved in the threat log.

- **Drop Packets**—Keeps all packets from continuing past the firewall.
- **Reset Both**—Resets the client and server.
- **Reset Client**—Resets the client.
- **Reset Server**—Resets the server.
- **Block-IP**—This action blocks traffic from either a source or a source-destination pair (configurable) for a specified period of time.

# Antivirus Profiles

▶ *Objects > Security Profiles > Antivirus*

A security policy can include specification of an antivirus profile to identify which applications are inspected for viruses and the action taken when a virus is detected. The profile can also specify actions to block spyware downloads. The default profile inspects all of the listed protocol decoders for viruses, generates alerts for Simple Mail Transport Protocol (SMTP), Internet Message Access Protocol (IMAP), and Post Office Protocol Version 3 (POP3), and takes the default action for other applications (alert or deny), depending on the type of virus detected.

Customized profiles can be used to minimize antivirus inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms.

For information on action types, refer to "Security Profiles" on page 211.

**Table 87. Antivirus Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of antivirus profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, periods, and underscores. |
| Description | Enter a description for the profile (up to 255 characters). |
| **Antivirus Tab** | |
| Packet Capture | Select the check box if you want to capture identified packets. |
| Decoders and Actions | For each type of traffic that you want to inspect for viruses, select an action from the drop-down list. You can also take specific action based on signatures created by WildFire. For more information on WildFire, refer to "About WildFire" on page 421. |

**Table 87. Antivirus Profile Settings (Continued)**

| Field | Description |
|---|---|
| Applications Exceptions and Actions | Identify applications that will be exceptions to the antivirus rule.<br>For example, to block all HTTP traffic except for a specific application, you can define an antivirus profile for which the application is an exception. **Block** is the action for the HTTP decoder, and **Allow** is the exception for the application.<br>To find an application, start typing the application name in the text box. A matching list of applications is displayed, and you can make a selection. The application is added to the table, and you can assign an action.<br>For each application exception, select the action to be taken when the threat is detected. |
| **Virus Exception Tab** | |
| Threat ID | Use this tab if you want the system to ignore specific threats. Exceptions that are already specified are listed. You can add additional threats by entering the threat ID and clicking **Add**. Threat IDs are presented as part of the threat log information. Refer to "Viewing the Logs" on page 264. |

# Anti-spyware Profiles

▶ *Objects > Security Profiles > Anti-spyware*

A security policy can include specification of an anti-spyware profile for "phone home" detection (detection of traffic from installed spyware). The default anti-spyware profile detects phone-home protection for all severity levels except the low and informational levels.

Customized profiles can be used to minimize anti-spyware inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms.

The Exceptions settings allows you to change the response to a specific signature. For example, you can block all packets that match a signature, except for the selected one, which generates an alert.

The **DNS Signatures** settings provides an additional method of identifying infected hosts on a network. These signatures detect specific DNS lookups for host names that have been associated with malware. The DNS signatures can be configured to allow, alert, or (default) block when these queries are observed, just as with regular antivirus signatures. Additionally, hosts that perform DNS queries for malware domains will appear in the botnet report. DNS signatures are downloaded as part of the antivirus updates.

The **Anti-spyware** page presents a default set of columns. Additional columns of information are available by using the column chooser. Click the arrow to the right of a column header and select the columns from the Columns sub-menu. For more information, refer to "Using Tables on Configuration Pages" on page 26.

**Table 88. Anti-spyware Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of anti-spyware profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, periods, and underscores. |
| Description | Enter a description for the profile (up to 255 characters). |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| **Rules Tab** | |
| Rule Name | Specify the rule name. |
| Threat Name | Enter **any** to match all signatures, or enter text to match any signature containing the entered text as part of the signature name. |
| Severity | Choose a severity level (**critical**, **high**, **medium**, **low**, or **informational**). |
| Action | Choose an action (**Default**, **Alert**, **Allow**, or **Drop**) for each threat. |
| Packet Capture | Select the check box if you want to capture identified packets. |
| **Exceptions Tab** | |
| Exceptions | Select the **Enable** check box for each threat for which you want to assign an action, or select **All** to respond to all listed threats. The list depends on the selected host, category, and severity. If the list is empty, there are no threats for the current selections.<br>Use the IP Address Exemptions column to add IP address filters to a threat exception. If IP addresses are added to a threat exception, the threat exception action for that signature will only be taken over the rule's action if the signature is triggered by a session having either the source or destination IP matching an IP in the exception. You can add up to 100 IP addresses per signature. With this option, you do not have to create a new policy rule and new vulnerability profile to create an exception for a specific IP address. |
| **DNS Signature Tab** | |
| Infected host DNS Signature | Choose an action to be taken when DNS lookups are made to known malware sites (Alert, Allow, or default (Block)). |
| Packet Capture | Select the check box if you want to capture identified packets. |
| Threat ID | Manually enter DNS signature exceptions (range 4000000-4999999). |

# Vulnerability Protection Profiles

▶ *Objects > Security Profiles > Vulnerability Protection*

A security policy can include specification of a vulnerability protection profile that determines the level of protection against buffer overflows, illegal code execution, and other attempts to exploit system vulnerabilities. The default profile protects clients and servers from all known critical, high, and medium-severity threats.

Customized profiles can be used to minimize vulnerability checking for traffic between trusted security zones, and to maximize protection for traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms. To apply vulnerability protection profiles to security policies, refer to "Security Policies" on page 187.

The Rules settings specify collections of signatures to enable, as well as actions to be taken when a signature within a collection is triggered.

The Exceptions settings allows you to change the response to a specific signature. For example, you can block all packets that match a signature, except for the selected one, which generates an alert. The **Exception** tab supports filtering functions.

The **Vulnerability Protection** page presents a default set of columns. Additional columns of information are available by using the column chooser. Click the arrow to the right of a column header and select the columns from the Columns sub-menu. For more information, refer to "Using Tables on Configuration Pages" on page 26.

**Table 89. Vulnerability Protection Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of vulnerability protection profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, periods, and underscores. |
| Description | Enter a description for the profile (up to 255 characters). |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| **Rules Tab** | |
| Rule Name | Specify a name to identify the rule. |
| Threat Name | Specify a text string to match. The firewall applies a collection of signatures to the rule by searching signature names for this text string. |
| Action | Choose the action (**Alert**, **Allow**, **Default**, or **Block**) to take when the rule is triggered. The **Default** action is based on the pre-defined action that is part of each signature provided by Palo Alto Networks. To view the default action for a signature, navigate to Objects > Security Profiles > Vulnerability Protection and click **Add** or select an existing profile. Click the **Exceptions** tab and then click **Show all signatures**. A list of all signatures will displayed and you will see an **Action** column. |
| Host | Specify whether to limit the signatures for the rule to those that are client side, server side, or either (**any**). |
| Packet Capture | Select the check box if you want to capture the packet that triggered the rule. |

**Table 89. Vulnerability Protection Profile Settings (Continued)**

| Field | Description |
|---|---|
| Category | Select a vulnerability category if you want to limit the signatures to those that match that category. |
| CVE List | Specify common vulnerabilities and exposures (CVEs) if you want to limit the signatures to those that also match the specified CVEs.<br>Each CVE is in the format CVE-*yyyy*-*xxxx*, where *yyyy* is the year and *xxxx* is the unique identifier. You can perform a string match on this field. For example, to find vulnerabilities for the year 2011, enter "2011". |
| Vendor ID | Specify vendor IDs if you want to limit the signatures to those that also match the specified vendor IDs.<br>For example, the Microsoft vendor IDs are in the form MS*yy*-*xxx*, where *yy* is the two-digit year and *xxx* is the unique identifier. For example, to match Microsoft for the year 2009, enter "MS09". |
| Severity | Select severities to match (**informational**, **low**, **medium**, **high**, or **critical**) if you want to limit the signatures to those that also match the specified severities. |
| **Exceptions Tab** | |
| Threats | Select the **Enable** check box for each threat for which you want to assign an action, or select **All** to respond to all listed threats. The list depends on the selected host, category, and severity. If the list is empty, there are no threats for the current selections.<br>Choose an action from the drop-down list box, or choose from the **Action** drop-down at the top of the list to apply the same action to all threats. If the **Show All** check box is selected, all signatures are listed. If the **Show All** check box is not selected, only the signatures that are exceptions are listed.<br>Select the **Packet Capture** check box if you want to capture identified packets.<br>The vulnerability signature database contains signatures that indicate a brute force attack; for example, Threat ID 40001 triggers on an FTP brute force attack. Brute-force signatures trigger when a condition occurs in a certain time threshold. The thresholds are pre-configured for brute force signatures, and can be changed by clicking the pencil icon next to the threat name on the **Vulnerability** tab (with the **Custom** option selected). You can specify the number of hits per unit of time and whether the threshold applies to source, destination, or source-and-destination.<br>Thresholds can be applied on a source IP, destination IP or a combination of source IP and destination IP.<br>***Note:*** *The default action is shown in parentheses. The* ***CVE*** *column shows identifiers for common vulnerabilities and exposures (CVE). These unique, common identifiers are for publicly known information security vulnerabilities.*<br>Use the IP Address Exemptions column to add IP address filters to a threat exception. If IP addresses are added to a threat exception, the threat exception action for that signature will only be taken over the rule's action if the signature is triggered by a session having either the source or destination IP matching an IP in the exception. You can add up to 100 IP addresses per signature. With this option, you do not have to create a new policy rule and new vulnerability profile to create an exception for a specific IP address. |

# URL Filtering Profiles

▶ *Objects > Security Profiles > URL Filtering*

A security policy can include specification of a URL filtering profile that blocks access to specific web sites and web site categories, or generates an alert when the specified web sites are accessed (a URL filtering license is required). You can also define a "block list" of web sites that are always blocked (or generate alerts) and an "allow list" of web sites that are always allowed. The web categories are predefined by Palo Alto Networks.

To apply URL filtering profiles to security policies, refer to "Security Policies" on page 187. To create custom URL categories with your own lists of URLs, refer to "Custom URL Categories" on page 242.

**Table 90. URL Filtering Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of URL filtering profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter a description for the profile (up to 255 characters). |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| Action on License Expiration | Select the action to take if the URL filtering license expires:<br>• **Block**—Blocks access to all web sites in the block list or the selected categories.<br>• **Allow**—Allows access to all web sites. |
| Dynamic URL Filtering | Select to enable dynamic URL categorization.<br>URL categorization takes advantage of a URL filtering database on the firewall that lists the most popular URLs and other URLs for malicious categories. The URL filtering database may be able to resolve requests that the local database is unable to categorize. The default is enabled when using the BrightCloud database. When using the PAN-DB, this option is enabled by default and is not configurable.<br>To configure the system response when a URL remains unresolved after a 5 second timeout period, use the Category and Action settings in this window (see Category Action later in this table). Select the action for the category "Not resolved URL." |
| Log Container Page Only | Select the check box to log only the URLs that match the content type that is specified. The default is enabled. |

**Table 90. URL Filtering Profile Settings (Continued)**

| Field | Description |
|---|---|
| Block List | Enter the IP addresses or URL path names of the web sites that you want to block or generate alerts on. Enter each URL one per line.<br><br>**IMPORTANT:** You must omit the "http and https" portion of the URLs when adding web sites to the list.<br>Entries in the block list are an exact match and are case-insensitive. For example, "www.paloaltonetworks.com" is different from "paloaltonetworks.com". If you want to block the entire domain, you should include both "*.paloaltonetworks.com" and "paloaltonetworks.com".<br>Examples:<br>• www.paloaltonetworks.com<br>• 198.133.219.25/en/US<br>Block and allow lists support wildcard patterns. The following characters are considered separators:<br>.<br>/<br>?<br>&<br>=<br>;<br>+<br>Every substring that is separated by the characters listed above is considered a token. A token can be any number of ASCII characters that does not contain any separator character or *. For example, the following patterns are valid:<br>*.yahoo.com (Tokens are: "*", "yahoo" and "com")<br>www.*.com (Tokens are: "www", "*" and "com")<br>www.yahoo.com/search=* (Tokens are: "www", "yahoo", "com", "search", "*")<br>The following patterns are invalid because the character "*" is not the only character in the token.<br>ww*.yahoo.com<br>www.y*.com |
| Action | Select the action to take when a web site in the block list is accessed.<br>• **alert**—Allow the user to access the web site, but add an alert to the URL log.<br>• **block**—Block access to the web site.<br>• **continue**—Allow the user to access the blocked page by clicking **Continue** on the block page.<br>• **override**—Allow the user to access the blocked page after entering a password. The password and other override settings are specified in the URL Admin Override area of the **Settings** page. Refer to Table 1 in the "Defining Management Settings" on page 30. |

**Table 90. URL Filtering Profile Settings (Continued)**

| Field | Description |
|---|---|
| Allow List | Enter the IP addresses or URL path names of the web sites that you want to allow or generate alerts on. Enter each URL one per line.<br><br>**IMPORTANT:** You must omit the "http and https" portion of the URLs when adding web sites to the list.<br>Entries in the allow list are an exact match and are case-insensitive. For example, "www.paloaltonetworks.com" is different from "paloaltonetworks.com". If you want to allow the entire domain, you should include both "*.paloaltonetworks.com" and "paloaltonetworks.com".<br>Examples:<br>• www.paloaltonetworks.com<br>• 198.133.219.25/en/US<br>Block and allow lists support wildcard patterns. The following characters are considered separators:<br>.<br>/<br>?<br>&<br>=<br>;<br>+<br>Every substring that is separated by the characters listed above is considered a token. A token can be any number of ASCII characters that does not contain any separator character or *. For example, the following patterns are valid:<br>*.yahoo.com (Tokens are: "*", "yahoo" and "com")<br>www.*.com (Tokens are: "www", "*" and "com")<br>www.yahoo.com/search=* (Tokens are: "www", "yahoo", "com", "search", "*")<br>The following patterns are invalid because the character "*" is not the only character in the token.<br>ww*.yahoo.com<br>www.y*.com<br>This list takes precedence over the selected web site categories. |
| Category/Action | For each category, select the action to take when a web site of that category is accessed.<br>• **alert**—Allow the user to access the web site, but add an alert to the URL log.<br>• **allow**—Allow the user to access the web site.<br>• **block**—Block access to the web site.<br>• **continue**—Allow the user to access the blocked page by clicking **Continue** on the block page.<br>• **override**—Allow the user to access the blocked page after entering a password. The password and other override settings are specified in the URL Admin Override area of the **Settings** page. Refer to Table 1 in the "Defining Management Settings" on page 30. |
| Check URL Category | Click to access the web site where you can enter a URL or IP address to view categorization information. |

# File Blocking Profiles

▶ *Objects > Security Profiles > File Blocking*

A security policy can include specification of a file blocking profile that blocks selected file types from being uploaded and/or downloaded, or generates an alert when the specified file types are detected. Table 92 lists the supported file formats at the time of this publication. However, because new file type support can be added in a content update, for the most up-to-date list, click **Add** in the **File Types** field of the File Blocking Profile dialog.

To apply file blocking profiles to security policies, refer to "Security Policies" on page 187.

**Table 91. File Blocking Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of file blocking profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter a description for the profile (up to 255 characters). |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| Rules | Define one or more rules to specify the action taken (if any) for the selected file types. To add a rule, specify the following and click **Add**:<br>• **Name**—Enter a rule name (up to 31 characters).<br>• **Applications**—Select the applications the rule applies to or select **any**.<br>• **File Types**—Select the file types for which you want to block or generate alerts.<br>• **Direction**—Select the direction of the file transfer (Upload, Download, or Both).<br>• **Action**—Select the action taken when the selected file types are detected:<br>– **alert**—An entry is added to the threat log.<br>– **block**—The file is blocked.<br>– **continue**—A message to the user indicates that a download has been requested and asks the user to confirm whether to continue. The purpose is to warn the user of a possible unknown download (also known as a drive-by-download) and to give the user the option of continuing or stopping the download.<br>***Note:*** *When you create a file blocking profile with the action* ***continue*** *or* ***continue-and-forward*** *(used for WildFire forwarding), you can only choose the application* ***web-browsing****. If you choose any other application, traffic that matches the security policy will not flow through the firewall due to the fact that the users will not be prompted with a continue page.*<br>– **forward**—The file is automatically sent to WildFire.<br>– **continue-and-forward**—A continue page is presented, and the file is sent to WildFire (combines the **continue** and **forward** actions). |

**Table 92. Supported File Formats for File Blocking**

| Field | Description |
|---|---|
| apk | Android application package file |
| avi | Video file based on Microsoft AVI (RIFF) file format |
| avi-divx | AVI video file encoded with the DivX codec |
| avi-xvid | AVI video file encoded with the XviD codec |
| bat | MS DOS Batch file |
| bmp-upload | Bitmap image file (upload only) |
| cab | Microsoft Windows Cabinet archive file |
| cdr | Corel Draw file |
| class | Java bytecode file |
| cmd | Microsoft command file |
| dll | Microsoft Windows Dynamic Link Library |
| doc | Microsoft Office Document |
| docx | Microsoft Office 2007 Document |
| dpx | Digital Picture Exchange file |
| dsn | Database Source Name file |
| dwf | Autodesk Design Web Format file |
| dwg | Autodesk AutoCAD file |
| edif | Electronic Design Interchange Format file |
| encrypted-doc | Encrypted Microsoft Office Document |
| encrypted-docx | Encrypted Microsoft Office 2007 Document |
| encrypted-office2007 | Encrypted Microsoft Office 2007 Fil |
| encrypted-pdf | Encrypted Adobe PDF Document |
| encrypted-ppt | Encrypted Microsoft Office PowerPoint |
| encrypted-pptx | Encrypted Microsoft Office 2007 PowerPoint |
| encrypted-rar | Encrypted rar file |
| encrypted-xls | Encrypted Microsoft Office Excel |
| encrypted-xlsx | Encrypted Microsoft Office 2007 Excel |
| encrypted-zip | Encrypted zip file |
| exe | Microsoft Windows Executable |
| flv | Adobe Flash Video file |
| gds | Graphics Data System file |
| gif-upload | GIF image file (upload only) |
| gzip | Files compressed with gzip utility |
| hta | HTML Application file |
| iso | Disc Image file based on ISO-9660 standard |

**Table 92. Supported File Formats for File Blocking (Continued)**

| Field | Description |
|---|---|
| iwork-keynote | Apple iWork Keynote documents |
| iwork-numbers | Apple iWork Numbers documents |
| iwork-pages | Apple iWork Pages documents |
| jar | Java ARchive |
| jpeg-upload | JPG/JPEG image file (upload only) |
| lnk | Microsoft Windows file shortcut |
| lzh | File compressed with lha/lzh utility/algorithm |
| mdb | Microsoft Access Database file |
| mdi | Microsoft Document Imaging file |
| mkv | Matroska Video file |
| mov | Apple Quicktime Movie file |
| mp3 | MP3 audio file |
| mp4 | MP4 audio file |
| mpeg | Movie file using MPEG-1 or MPEG-2 compression |
| msi | Microsoft Windows Installer package file |
| msoffice | Microsoft Office File (doc, xls, ppt, pub, pst) |
| ocx | Microsoft ActiveX file |
| pdf | Adobe Portable Document file |
| PE | Microsoft Windows Portable Executable (exe, dll, com, scr, ocx, cpl, sys, drv, tlb) |
| pgp | Security key or digital signature encrypted with PGP software |
| pif | Windows Program Information File containing executable instructions |
| pl | Perl Script file |
| png-upload | PNG image file (upload only) |
| ppt | Microsoft Office PowerPoint Presentation |
| pptx | Microsoft Office 2007 PowerPoint Presentation |
| psd | Adobe Photoshop Document |
| rar | Compressed file created with winrar |
| reg | Windows Registry file |
| rm | RealNetworks Real Media file |
| rtf | Windows Rich Text Format document file |
| sh | Unix Shell Script file |
| stp | STandard for the Exchange of Product model data 3D graphic file |
| tar | Unix tar archive file |
| tdb | Tanner Database (www.tannereda.com) |
| tif | Windows Tagged Image file |

**Table 92. Supported File Formats for File Blocking (Continued)**

| Field | Description |
|---|---|
| torrent | BitTorrent file |
| wmf | Windows Metafile to store vector images |
| wmv | Windows Media Video file |
| wri | Windows Write document file |
| wsf | Windows Script file |
| xls | Microsoft Office Excel |
| xlsx | Microsoft Office 2007 Excel |
| zcompressed | Compressed Z file in Unix, decompressed with uncompress |
| zip | Winzip/pkzip file |

# Data Filtering Profiles

▶ *Objects > Security Profiles > Data Filtering*

A security policy can include specification of a data filtering profile to help identify sensitive information such as credit card or social security numbers and prevent the sensitive information from leaving the area protected by the firewall.

To apply data filtering profiles to security policies, refer to "Security Policies" on page 187.

**Table 93. Data Filtering Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of log forwarding profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter a description for the profile (up to 255 characters). |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| Data Capture | Select the check box to automatically collect the data that is blocked by the filter. |

***Note:*** *Specify a password for Manage Data Protection on the* ***Settings*** *page to view your captured data. Refer to "Defining Management Settings" on page 30.*

To add a data pattern, click **Add** and specify the following information.

**Table 94. Data Pattern Settings**

| Field | Description |
|---|---|
| Data Pattern | Choose an existing data pattern from the Data Pattern drop-down list, or configure a new pattern by choosing **Data Pattern** from the list and specifying the following information:<br>• **Name**—Configure a name for the data pattern.<br>• **Description**—Configure a description for the data pattern.<br>• **Shared**—Select this option if the data pattern object will be shared across multiple virtual systems.<br>• **Weight**—Specify unit values for the specified patterns to use in calculating thresholds. For instance, if you designate a weight of 5 for SSN#, every instance of a SSN pattern will increment the threshold by 5. In other words, the detection of ten SSN patterns will result in 10 x 5 (weight) = 50.<br>– **CC#**—Specify a weight for the credit card field (range 0-255).<br>– **SSN#**—Specify a weight for the social security number field, where the field includes dashes, such as 123-45-6789 (range 0-255, 255 is highest weight).<br>– **SSN# (without dash)**—Specify a weight for the social security number field, where the entry is made without dashes, such as 123456789 (range 0-255, 255 is highest weight).<br>• **Custom Patterns**—To match a custom data pattern for the traffic that is subject to this profile, create a custom data pattern by clicking **Add** and specifying the pattern name, regular expression (regex) to match, and weight (0-255, 255 is highest weight). You can add multiple match expressions to the same data pattern profile. |
| Applications | Specify the applications to include in the filtering rule:<br>• Choose **any** to apply the filter to all of the listed applications. This selection does not block all possible applications, just the listed ones.<br>• Click **Add** to specify individual applications. |
| File Types | Specify the file types to include in the filtering rule:<br>• Choose **any** to apply the filter to all of the listed file types. This selection does not block all possible file types, just the listed ones.<br>• Click **Add** to specify individual file types. |
| Direction | Specify whether to apply the filter in the upload direction, download direction, or both. |
| Alert Threshold | Specify the value that will trigger an alert. For example, if you have a threshold of 100 with a SSN weight of 5, the rule will need to detect at least 20 SSN patterns before the rule will be triggered (20 instances x 5 weight = 100). |
| Block Threshold | Specify the value that will trigger a block. For example, if you have a threshold of 100 with a SSN weight of 5, the rule will need to detect at least 20 SSN patterns before the rule will be triggered (20 instances x 5 weight = 100). |

# DoS Profiles

▶ *Objects > Security Profiles > DoS Protection*

A DoS protection policy can include specification of a DoS profile to protect against DoS attacks and take protective action in response to rule matches. The DoS profile specifies the types of actions and the matching criteria.

To apply DoS profiles to DoS policies, refer to "DoS Protection Policies" on page 209.

**Table 95. DoS Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of log forwarding profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| Description | Enter a description of the profile (up to 255 characters). |
| Type | Specify one of the following profile types:<br>• **aggregate**—Apply the DoS thresholds configured in the profile to all packets that match the rule criteria on which this profile is applied. For example, an aggregate rule with a SYN flood threshold of 10000 packets per second (pps) counts all packets that hit that particular DoS rule.<br>• **classified**—Apply the DoS thresholds configured in the profile to all packets satisfying the classification criterion (source IP, destination IP or source-and-destination IP). |
| **Flood Protection Tab** | |
| Syn Flood subtab<br>UDP Flood subtab<br>ICMP Flood subtab<br>Other subtab | Select the check box to enable SYN flood protection, and specify the following settings:<br>• **Choice**—(SYN Flood only) Choose from the following options:<br>– **Random early drop**—Drop packets randomly before the overall DoS limit is reached.<br>– **SYN cookies**—Use SYN cookies to generate acknowledgments so that it is not necessary to drop connections in the presence of a SYN flood attack.<br>• **Alarm Rate**—Specify the rate (pps) at which a DoS alarm is generated (range 0-2000000 pps, default 10000 pps).<br>• **Activate Rate**—Specify the rate (pps) at which a DoS response is activated (range 0-2000000 pps, default 10000 pps).<br>• **Maximal Rate**—Specify the rate at which packets will be dropped or blocked.<br>• **Block Duration**—Specify the length of time (seconds) during which the offending packets will be denied. Packets arriving during the block duration do not count towards triggered alerts.<br>***Note:*** *When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session.* |

**Table 95. DoS Profile Settings (Continued)**

| Field | Description |
|---|---|
| **Resources Protection Tab** | |
| Sessions | Select the check box to enable resources protection. |
| Max Concurrent Limit | Specify the maximum number of concurrent sessions. If the DoS profile type is aggregate, this limit applies to the entire traffic hitting the DoS rule on which the DoS profile is applied. If the DoS profile type is classified, this limit applies to the entire traffic on a classified basis (source IP, destination IP or source-and-destination IP) hitting the DoS rule on which the DoS profile is applied. |

# Other Policy Objects

Policy objects are the elements that enable you to construct, schedule, and search for policies. The following element types are supported:

- Addresses and address groups to determine the scope of the policy. Refer to "Addresses and Address Groups" in the next section.
- Applications and application groups that allow you to specify how software applications are treated in policies. Refer to "Applications and Application Groups" on page 231.
- Application filters that allow you to simplify searches. Refer to "Application Filters" on page 238.
- Services and service groups to limit the port numbers. Refer to "Services" on page 239.
- Data patterns to define categories of sensitive information for data filtering policies. Refer to "Data Patterns" on page 240.
- Custom URL categories that contain your own lists of URLs to include as a group in URL filtering profiles. Refer to "Custom URL Categories" on page 242.
- Spyware and vulnerability threats to allow for detailed threat responses. Refer to "Security Profile Groups" on page 246.
- Log forwarding to specify log settings. Refer to "Log Forwarding" on page 247.
- Schedules to specify when policies are active. Refer to "Schedules" on page 250.

# Addresses and Address Groups

To define security policies for specific source or destination addresses, you must first define the addresses and address ranges. Addresses requiring the same security settings can be combined into address groups that you can refer to as a unit.

## Defining Address Ranges

▶ *Objects > Addresses*

To define security policies for specific source or destination addresses, you must first define the addresses and address ranges. Addresses requiring the same security settings can be combined into address groups to simplify policy creation (refer to "Defining Address Groups" on page 230).

**Table 96. New Address Settings**

| Field | Description |
|---|---|
| Name | Enter a name that describes the addresses to be defined (up to 63 characters). This name appears in the address list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow use by all virtual systems. |
| Description | Enter a description for the object (up to 255 characters). |
| Type | Specify an IPv4 or IPv6 address or address range, FQDN, or Dynamic.<br>**IP Netmask:**<br>Enter the IPv4 or IPv6 address or IP address range using the following notation:<br>*ip_address/mask* or *ip_address*<br>where the *mask* is the number of significant binary digits used for the network portion of the address.<br>Example:<br>"192.168.80.150/32" indicates one address, and "192.168.80.0/24" indicates all addresses from 192.168.80.0 through 192.168.80.255.<br>Example:<br>"2001:db8:123:1::1" or "2001:db8:123:1::/64"<br>**IP Range:**<br>To specify an address range, select **IP Range**, and enter a range of addresses. The format is:<br>*ip_address–ip_address*<br>where each address can be IPv4 or IPv6.<br>Example:<br>"2001:db8:123:1::1 - 2001:db8:123:1::22" |

**Table 96. New Address Settings (Continued)**

| Field | Description |
|---|---|
| Type (continued) | **FQDN:**<br>To specify an address using the FQDN, select **FQDN** and enter the domain name.<br>The FQDN initially resolves at commit time. Entries are subsequently refreshed when the firewall performs a check every 30 minutes; all changes in the IP address for the entries are picked up at the refresh cycle.<br>The FQDN is resolved by the system DNS server or a DNS proxy object, if a proxy is configured. For information about DNS proxy, refer to "DNS Proxy" on page 173.<br>**Dynamic:**<br>Specify a name, description, and an identifier that will be used by the XML API to feed in addresses. The identifier can be any string, must be unique and is not limited to a UUID standard. Once you have the name and identifier specified, you will use the "name" of the object in your policies and the identifier (the field located to the right of **"Type:** Dynamic" is used to map the unique IP address to an address object in the XMP API script.<br> 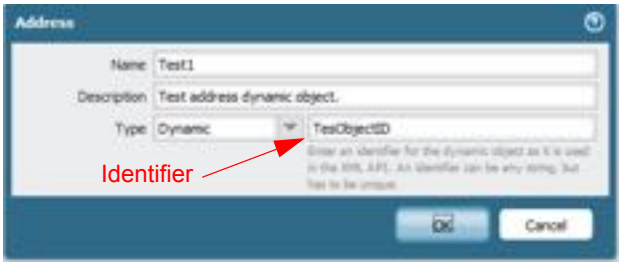 <br>This dynamic address object feature is useful in environments where host IP addresses change frequently; like in virtual environments where hosts are often added and removed and where failovers may occur requiring a change in IP address. By using a name for the dynamic address object for policy matching instead of static IP addresses or networks, the addresses can be dynamically updated without modifying the firewall configuration.<br>Each Dynamic Address Object can contain up to 256 IP addresses. |

## Defining Address Groups

▶ *Objects > Address Groups*

To simplify the creation of security policies, addresses requiring the same security settings can be combined into address groups.

**Table 97. Address Group**

| Field | Description |
|---|---|
| Name | Enter a name that describes the address group (up to 63 characters). This name appears in the address list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow use by all virtual systems. |
| Addresses | Click Add and select addresses and/or other address groups to be included in this group. |

## Defining Regions

▶ *Objects > Regions*

The firewall supports creation of policy rules that apply to specified countries or other regions. The region is available as an option when specifying source and destination for security policies, decryption policies, and DoS policies. You can choose from a standard list of countries or use the region settings described in this section to define custom regions to include as options for security policy rules.

**Table 98. New Address Settings**

| Field | Description |
|---|---|
| Name | Enter a name that describes the region (up to 31 characters). This name appears in the address list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Geo Location | To specify latitude and longitude, select the check box and values (*xxx.xxxxxx* format). This information is used in the traffic and threat maps for App-Scope. Refer to "Using App-Scope" on page 256. |
| Addresses | Specify an IP address, range of IP addresses, or subnet to identify the region, using any of the following formats:<br>*x.x.x.x*<br>*x.x.x.x-y.y.y.y*<br>*x.x.x.x/n* |

# Applications and Application Groups

▶ *Objects > Applications*

The **Applications** page lists various attributes of each application definition, such as the application's relative security risk (1 to 5). The risk value is based on criteria such as whether the application can share files, is prone to misuse, or tries to evade firewalls. Higher values indicate higher risk.

The top application browser area of the page lists the attributes that you can use to filter the display. The number to the left of each entry represents the total number of applications with that attribute.

You can perform any of the following functions on this page:

- To apply application filters, click an item that you want to use as a basis for filtering. For example, to restrict the list to the Networking category, click **Networking** and the list will only show networking applications.

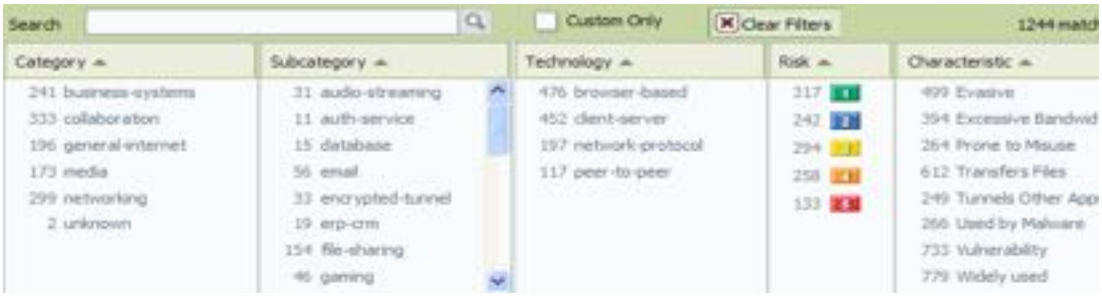


- To filter on additional columns, select an entry in the other columns. The filtering is successive: first category filters are applied, then subcategory filters, then technology filters, then risk, and finally characteristic filters.

  For example, the next figure shows the result of applying a category, subcategory, and risk filter. In applying the first two filters, the **Technology** column is automatically restricted to the technologies that are consistent with the selected category and sub category, even though a technology filter has not been explicitly applied.

  Each time a filter is applied, the list of applications in the lower part of the page is automatically updated, as shown in the following figure. Any saved filters can be viewed in **Objects > Application Filters**.

- To search for a specific application, enter the application name or description in the **Search** field, and press **Enter**. The application is listed, and the filter columns are updated to show statistics for the applications that matched the search.

  A search will match partial strings. When you define security policies, you can write rules that apply to all applications that match a saved filter. Such rules are dynamically updated when a new application is added through a content update that matches the filter.

- Click an application name to view additional details about the application, as described in the following table. You can also customize risk and timeout values, as described in the following table.

**Table 99. Application Details**

| Item | Description |
|---|---|
| Name | Name of the application. |
| Description | Description of the application (up to 255 characters). |
| Additional Information | Links to web sources (Wikipedia, Google, and Yahoo!) that contain additional information about the application. |
| Standard Ports | Ports that the application uses to communicate with the network. |
| Capable of File Transfer | Indication of whether the application is able to transfer files. |
| Used by Malware | Indication of whether the application is used by malware. |
| Excessive Bandwidth Use | Indication of whether the application uses too much bandwidth so that network performance may be compromise. |
| Evasive | Indication of whether the application attempts to evade firewalls. |
| Widely used | Indication of whether the effects of the application are wide-ranging. |
| Has Known Vulnerabilities | Indication of whether the application has any currently known vulnerabilities. |
| Tunnels Other Applications | Indication of whether the application can carry other applications within the messages that it sends. |
| Depends on Applications | List of other applications that are required for this application to run. |
| Category | Application category. |
| Subcategory | Application sub category. |
| Technology | Application technology. |
| Risk | Assigned risk of the application.<br>To customize this setting, click the **Customize** link, enter a value (1-5), and click **OK**. |
| Prone to Misuse | Indication of whether the application tends to attract misuse. |
| Session Timeout | Period of time (seconds) required for the application to timeout due to inactivity (1-604800 seconds). This timeout is for protocols other than TCP or UDP. For TCP and UDP, refer to the next rows in this table.<br>To customize this setting, click the **Customize** link, enter a value (seconds), and click **OK**. |

**Table 99. Application Details (Continued)**

| Item | Description |
|---|---|
| TCP Timeout (seconds) | Timeout for terminating a TCP application flow (1-604800 seconds).<br>To customize this setting, click the **Customize** link, enter a value (seconds), and click **OK**. |
| UDP Timeout (seconds): | Timeout for terminating a UDP application flow (1-604800 seconds).<br>To customize this setting, click the **Customize** link, enter a value (seconds), and click **OK**. |

When the firewall is not able to identify an application using the application ID, the traffic is classified as unknown: unknown-tcp or unknown-udp. This behavior applies to all unknown applications except those that fully emulate HTTP. For more information, refer to "Identifying Unknown Applications and Taking Action" on page 276.

You can create new definitions for unknown applications and then define security policies for the new application definitions. In addition, applications that require the same security settings can be combined into application groups to simplify the creation of security policies.

## Defining Applications

▶ *Objects > Applications*

Use the **Applications** page to add new applications for the firewall to evaluate when applying policies.

**Table 100. New Application Settings**

| Field | Description |
|---|---|
| **Configuration Tab** | |
| Name | Enter the application name (up to 31 characters). This name appears in the applications list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, periods, hyphens, and underscores. The first character must be a letter. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the application to be shared by all virtual systems. |
| Description | Enter a description of the application for general reference (up to 255 characters). |
| Category | Select the application category, such as email or database. For a description of each category, refer to "Application Categories and Subcategories" on page 437. The category is used to generate the Top Ten Application Categories chart and is available for filtering (refer to "Using the Application Command Center" on page 253). |
| Subcategory | Select the application subcategory, such as email or database. For a description of each sub category, refer to "Application Categories and Subcategories" on page 437. The sub category is used to generate the Top Ten Application Categories chart and is available for filtering (refer to "Using the Application Command Center" on page 253). |
| Technology | Select the technology for the application. For a description of each technology, refer to "Application Technologies" on page 439. |

**Table 100. New Application Settings (Continued)**

| Field | Description |
|---|---|
| Parent App | Specify a parent application for this application. This setting applies when a session matches both the parent and the custom applications; however, the custom application is reported because it is more specific. |
| Risk | Select the risk level associated with this application (1=lowest to 5=highest). |
| Characteristics | Select the application characteristics that may place the application at risk. For a description of each characteristic, refer to "Application Characteristics" on page 439. |
| **Advanced Tab** | |
| Port | If the protocol used by the application is TCP and/or UDP, select **Port** and enter one or more combinations of the protocol and port number (one entry per line). The general format is:<br><*protocol*>/<*port*><br>where the <*port*> is a single port number, or **dynamic** for dynamic port assignment.<br>Examples: TCP/dynamic or UDP/32.<br>This setting applies when using **app-default** in the **Service** column of a security rule. |
| IP Protocol | To specify an IP protocol other than TCP or UDP, select **IP Protocol**, and enter the protocol number (1 to 255). |
| ICMP Type | To specify an Internet Control Message Protocol version 4 (ICMP) type, select **ICMP Type** and enter the type number (range 0-255). |
| ICMP6 Type | To specify an Internet Control Message Protocol version 6 (ICMPv6) type, select **ICMP6 Type** and enter the type number (range 0-255). |
| None | To specify signatures independent of protocol, select **None**. |
| Timeout | Enter the number of seconds before an idle application flow is terminated (range 0-604800 seconds). A zero indicates that the default timeout of the application will be used. This value is used for protocols other than TCP and UDP in all cases and for TCP and UDP timeouts when the TCP timeout and UDP timeout are not specified. |
| TCP Timeout | Enter the number of seconds before an idle TCP application flow is terminated (range 0-604800 seconds). A zero indicates that the default timeout of the application will be used. |
| UDP Timeout | Enter the number of seconds before an idle UDP application flow is terminated (range 0-604800 seconds). A zero indicates that the default timeout of the application will be used. |
| Scanning | Select check boxes for the scanning types that you want to allow, based on security profiles (file types, data patterns, and viruses). |

**Table 100. New Application Settings (Continued)**

| Field | Description |
|---|---|
| **Signature Tab** | |
| Signatures | Click **Add** to add a new signature, and specify the following information:<br>• **Signature Name**—Enter a name to identify the signature.<br>• **Comment**—Enter an optional description.<br>• **Scope**—Select whether to apply this signature only to the current transaction or to the full user session.<br>• **Ordered Condition Match**—Select if the order in which signature conditions are defined is important.<br>Specify conditions to define signatures:<br>• Add a condition by clicking **Add AND Condition** or **Add OR Condition**. To add a condition within a group, select the group and then click **Add Condition**.<br>• Select an operator from **Pattern Match** and **Equal To**. When choosing a pattern match operator, specify the following:<br>– **Context**—Select from the available contexts.<br>– **Pattern**—Specify a regular expression. See Table 104 for pattern rules for regular expressions.<br>– **Qualifier and Value**—Optionally, add qualifier/value pairs.<br>• When choosing an equal to operator, specify the following,<br>– **Context**—Select from unknown requests and responses for TCP or UDP.<br>– **Position**—Select between the first four or second four bytes in the payload.<br>– **Mask**—Specify a 4-byte hex value, for example, 0xffffff00.<br>– **Value**—Specify a 4-byte hex value, for example, 0xaabbccdd.<br>• To move a condition within a group, select the condition and click the **Move Up** or **Move Down** arrow. To move a group, select the group and click the **Move Up** or **Move Down** arrow. You cannot move conditions from one group to another. |

***Note:*** *It is not required to specify signatures for the application if the application is used only for application override rules.*

To import an application, click **Import**. Browse to select the file, and select the target virtual system from the **Destination** drop-down list.

To export the application, select the check box for the application and click **Export**. Follow the prompts to save the file.

## Custom Applications with Signatures

You can define custom applications with signatures. This section provides examples of how this can be done. Refer to the *PAN-OS Command Line Interface Reference Guide* for information on the show application command.

### Example - Detect web traffic to a specified site

This example shows an application that detects web traffic going to *www.specifiedsite.com*.

Requests to the web site are of the following form:

```
GET /001/guest/
viewprofile.act?fa=25&tg=M&mg=F&searchType=zipcode&type=QUICK&pict=true&cont
ext=adrr&zip=94024&ta=34&sb=&item=0&pn=0 HTTP/1.1

Host: www.specifiedsite.com

User-Agent: Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US; rv:1.9.0.7)
Gecko/2009021910 Firefox/3.0.7 Accept: text/html,application/
xhtml+xml,application/xml;q=0.9,*/*;q=0.8 Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip,deflate Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
Keep-Alive: 300 Connection: keep-alive Referer: http://www.specifiedsite.com/
001/guest/
search.act?type=QUICK&pict=true&sb=&fa=25&ta=34&mg=F&tg=M&searchType=zipcode
&zip=94024&context=adrr&context=adrr Cookie:
JSESSIONID=A41B41A19B7533589D6E88190B7F0B3D.001; specifiedsite.com/
jumpcookie=445461346*google.com/search?q=lava+life&; locale=en_US;
campaign=1; imageNum=2; cfTag_LogSid=9327803497943a1237780204643;
__utma=69052556.1949878616336713500.1238193797.1238193797.1238193797.1;
__utmb=69052556.2.10.1238193797; __utmc=69052556;
__utmz=69052556.1238193797.1.1.utmcsr=(direct)|utmccn=(direct)|utmcmd=(none)
; __utmv=69052556.gender%3Df; launch=1
```

The following signature can identify *specifiedsite* traffic if the host field is *www.specifiedsite.com*.

```
username@hostname# show application specifiedsite

specifiedsite {
  category collaboration;
  subcategory social-networking;
  technology browser-based;
  decoder http;
  signature {
    s1 {
      and-condition {
        a1 {
          or-condition {
            o1 {
              context http-req-host-header;
              pattern www\.specifiedsite\.com;
            }
          }
        }
      }
    }
  }
}
```

### Example - Detect a post to a specified blog

This example shows an application that detects blog posting activity on *www.specifiedblog.com*. In this example, it is not necessary to detect when somebody tries to read the blog, only to detect when an item is getting posted.

The post traffic request includes the following:

```
POST /wp-admin/post.php HTTP/1.1 Host: panqa100.specifiedblog.com
User-Agent: Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US; rv:1.9.0.7)
```

```
Gecko/2009021910 Firefox/3.0.7 Accept: text/html,application/
xhtml+xml,application/xml;q=0.9,*/*;q=0.8 Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip,deflate Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
Keep-Alive: 300 Connection: keep-alive Referer: http://
panqa100.specifiedblog.com/wp-admin/post.php?action=edit&post=1
Cookie: __utma=96763468.235424814.1238195613.1238195613.1238195613.1;
__utmb=96731468; __utmc=96731468;
__utmz=96731468.1238195613.1.1.utmccn=(organic)|utmcsr=google|utmctr=blog+ho
st|utmcmd=organic; wordpressuser_bfbaae4493589d9f388265e737a177c8=panqa100;
wordpresspass_bfbaae4493589d9f388265e737a177c8=c68a8c4eca4899017c58668eacc05
fc2
Content-Type: application/x-www-form-urlencoded Content-Length: 462
user_ID=1&action=editpost&post_author=1&post_ID=1&post_title=Hello+world%21&
post_category%5B%5D=1&advanced_view=1&comment_status=open&post_password=&exc
erpt=&content=Hello+world.%3Cbr+%2F%3E&use_instant_preview=1&post_pingback=1
&prev_status=publish&submit=Save&referredby=http%3A%2F%2Fpanqa100.specifiedb
log.com%2Fwp-admin%2F&post_status=publish&trackback_url=&post_name=hello-
world&post_author_override=1&mm=3&jj=27&aa=2009&hh=23&mn=14&ss=42&metakeyinp
ut=&metavalue=HTTP/1.1
```

The host field includes the pattern *specifiedblog.com*. However, if a signature is written with that value in the host, it will match all traffic going to *specifiedblog*.com, including posting and viewing traffic. Therefore, it is necessary to look for more patterns.

One way to do this is to look for *post_title* and *post-author* patterns in the parameters of the post. The resulting signature detects postings to the web site:

```
username@hostname# show application specifiedblog_blog_posting
specifiedblog_blog_posting {
  category collaboration;
  subcategory web-posting;
  technology browser-based;
  decoder http;
  signature {
    s1 {
      and-condition {
        a1 {
          or-condition {
            o1 {
              context http-req-host-header;
              pattern specifiedblog\.com;
              method POST;
            }
          }
        }
        a2 {
          or-condition {
            o2 {
              context http-req-params;
              pattern post_title;
              method POST;
            }
          }
        }
        a3 {
          or-condition {
            o3 {
              context http-req-params;
              pattern post_author;
              method POST;
            }
          }
        }
      }
    }
  }
}
```

## Defining Application Groups

▶ *Objects > Application Groups*

To simplify the creation of security policies, applications requiring the same security settings can be combined into application groups. To define new applications, refer to "Defining Applications" on page 233.

**Table 101. New Application Group**

| Field | Description |
|---|---|
| Name | Enter a name that describes the application group (up to 31 characters). This name appears in the application list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Applications | Click **Add** and select applications, application filters, and/or other application groups to be included in this group. |

# Application Filters

▶ *Objects > Application Filters*

You can define application filters to simplify repeated searches. To define application filters to simplify repeated searches, click **Add** and enter a name for the filter.

In the upper area of the window, click an item that you want to use as a basis for filtering. For example, to restrict the list to the Networking category, click **networking**.

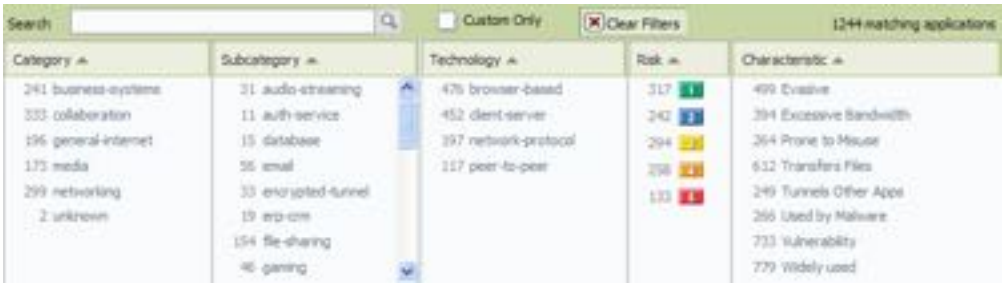


To filter on additional columns, select an entry in the columns to display check boxes. The filtering is successive: first category filters are applied, then sub category filters, then technology filters, then risk, filters, and finally characteristic filters.

For example, the next figure shows the result of choosing a category, sub category, and risk filter. In applying the first two filters, the **Technology** column is automatically restricted to the technologies that are consistent with the selected category and sub category, even though a technology filter has not been explicitly applied.

As you select options, the list of applications in the lower part of the page is automatically updated, as shown in the figure.

# Services

▶ *Objects > Services*

When you define security policies for specific applications, you can select one or more services to limit the port numbers the applications can use. The default service is **any**, which allows all TCP and UDP ports.

The HTTP and HTTPS services are predefined, but you can add additional service definitions. Services that are often assigned together can be combined into service groups to simplify the creation of security policies (refer to "Service Groups" on page 240).

**Table 102. Service Settings**

| Field | Description |
|---|---|
| Name | Enter the service name (up to 63 characters). This name appears in the services list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter a description for the service (up to 255 characters). |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems. |
| Protocol | Select the protocol used by the service (TCP or UDP). |
| Destination Port | Enter the destination port number (0 to 65535) or range of port numbers (port1-port2) used by the service. Multiple ports or ranges must be separated by commas. The destination port is required. |
| Source Port | Enter the source port number (0 to 65535) or range of port numbers (port1-port2) used by the service. Multiple ports or ranges must be separated by commas. The source port is optional. |

# Service Groups

▶ *Objects > Services Groups*

To simplify the creation of security policies, you can combine services that have the same security settings into service groups. To define new services, refer to "Services" on page 239.

**Table 103. Service Group Settings**

| Field | Description |
|---|---|
| Name | Enter the service group name (up to 63 characters). This name appears in the services list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Service | Click **Add** to add services to the group. Select from the drop-down list, or click the **Service** button at the bottom of the drop-down list, and specify the settings. Refer to "Services" on page 239 for a description of the settings. |

# Data Patterns

Data pattern support allows you to specify categories of sensitive information that you may want to subject to filtering using data filtering security policies. For instructions on configuring data patterns, refer to "Dynamic Block Lists" on page 242.

When adding a new pattern (regular expression), the following general requirements apply:

- The pattern must have string of at least 7 bytes to match. It can contain more than 7 bytes, but not fewer.
- The string match is case-sensitive, as with most regular expression engines. Looking for "confidential" is different than looking for "Confidential" or "CONFIDENTIAL."

The regular expression syntax in PAN-OS is similar to traditional regular expression engines, but every engine is unique. The following table describes the syntax supported in PAN-OS.

**Table 104. Pattern Rules**

| Syntax | Description |
|---|---|
| . | Match any single character. |
| ? | Match the preceding character or expression 0 or 1 time. The general expression MUST be inside a pair of parentheses.<br>Example: (abc)? |
| * | Match the preceding character or expression 0 or more times. The general expression MUST be inside a pair of parentheses.<br>Example: (abc)* |
| + | Match the preceding character or regular expression 1 or more times. The general expression MUST be inside a pair of parentheses.<br>Example: (abc)+ |

**Table 104. Pattern Rules**

| Syntax | Description |
|---|---|
| \| | Equivalent to "or".<br>Example: ((bif)\|(scr)\|(exe)) matches "bif", "scr" or "exe". Note that the alternative substrings must be in parentheses. |
| - | Used to create range expressions.<br>Example: [c-z] matches any character between c and z, inclusive. |
| [ ] | Match any.<br>Example: [abz]: matches any of the characters a, b, or z. |
| ^ | Match any except.<br>Example: [^abz] matches any character except a, b, or z. |
| { } | Min/Max number of bytes.<br>Example: {10-20} matches any string that is between 10 and 20 bytes. This must be directly in front of a fixed string, and only supports "-". |
| \ | To perform a literal match on any one of the special characters above, it MUST be escaped by preceding them with a '\' (backslash). |
| &amp | & is a special character, so to look for the "&" in a string you must use "&amp" instead. |

### Data Patterns Examples

The following are examples of valid custom patterns:

- .*((Confidential)|(CONFIDENTIAL))
  - Looks for the word "Confidential" or "CONFIDENTIAL" anywhere
  - ".*" at the beginning specifies to look anywhere in the stream
  - Does not match "confidential" (all lower case)
- .*((Proprietary &amp Confidential)|(Proprietary and Confidential))
  - Looks for either "Proprietary & Confidential" or "Proprietary and Confidential"
  - More precise than looking for "Confidential"
- .*(Press Release).*((Draft)|(DRAFT)|(draft))
  - Looks for "Press Release" followed by various forms of the word draft, which may indicate that the press release isn't ready to be sent outside the company
- .*(Trinidad)
  - Looks for a project code name, such as "Trinidad"

# Custom URL Categories

▶ *Objects > Custom URL Categories*

The custom URL categories feature allows you to create your own lists of URLs that can be selected in any URL filtering profile. Each custom category can be controlled independently and will have an action associated with it in each URL filtering profile (allow, block, continue, override, or alert).

URL entries can be added individually, or you can import a list of URLs. To do so, create a text file that contains the URLs to include, with one URL per line. Each URL can be in the format "www.example.com," and can contain * as a wildcard, such as "*.example.com." For additional information on wildcards, refer to the description of Block List in Table 90 on page 217.

***Note:*** *URL entries added to custom categories are case insensitive.*

For instructions on setting up URL filtering profiles, refer to "URL Filtering Profiles" on page 217.

**Table 105. Custom URL Categories**

| Field | Description |
|---|---|
| Name | Enter a name to identify the custom URL category (up to 31 characters). This name appears in the category list when defining URL filtering policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter a description for the URL category (up to 255 characters). |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| Sites | In the Sites area, click **Add** to enter a URL or click **Import** and browse to select the text file that contains the list of URLs. |

# Dynamic Block Lists

▶ *Objects > Dynamic Block Lists*

Use the **Dynamic Block Lists** page to create an address object based on an imported list of IP addresses. The source of the list must be a text file and must be located on a web server. You can set the **Repeat** option to automatically update the list on the device hourly, daily, weekly, or monthly. After creating a dynamic block list object, you can then use the address object in the source and destination fields for security policies. Each imported list can contain up to 5,000 IP addresses (IPv4 and/or IPv6), IP ranges, or subnets.

The list must contain one IP address, range, or subnet per line, for example:

"192.168.80.150/32" indicates one address, and "192.168.80.0/24" indicates all addresses from 192.168.80.0 through 192.168.80.255.

Example:

"2001:db8:123:1::1" or "2001:db8:123:1::/64"

**IP Range:**

To specify an address range, select **IP Range**, and enter a range of addresses. The format is:
*ip_address–ip_address*
where each address can be IPv4 or IPv6.
Example:
"2001:db8:123:1::1 - 2001:db8:123:1::22"

**Table 106 Dynamic Block Lists**

| Field | Description |
|---|---|
| Name | Enter a name to identify the Dynamic Block List (up to 32 characters). This name will appear when selecting the source or destination in a policy. |
| Description | Enter a description for the block list (up to 255 characters). |
| Source | Enter an HTTP or HTTPS URL path that contains the text file. For example, http://1.1.1.1/myfile.txt. |
| Repeat | Specify the frequency in which the list should be imported. You can choose hourly, daily, weekly, or monthly. At the specified interval, the list will be imported into the configuration. A full commit is not needed for this type of update to occur. |
| Test Source URL | Test that the source URL or server path is available. |

# Custom Spyware and Vulnerability Signatures

This section describes the options available to create custom Spyware and Vulnerability signatures that can be used when creating custom vulnerability profiles.

- ***Objects > Custom Signatures > Data Patterns***
- ***Objects > Custom Signatures > Spyware***
- ***Objects > Custom Signatures > Vulnerability***

## Defining Data Patterns

- ***Objects > Custom Signatures > Data Patterns***

Use the **Data Patterns** page to define the categories of sensitive information that you may want to subject to filtering using data filtering security policies. For information on defining data filtering profiles, refer to "Data Filtering Profiles" on page 223.

**Table 107. Data Pattern Settings**

| Field | Description |
|---|---|
| Name | Enter the data pattern name (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter a description for the data pattern (up to 255 characters). |

**Table 107. Data Pattern Settings (Continued)**

| Field | Description |
|---|---|
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| Weight | Enter weights for pre-specified pattern types. The weight is a number between 1 and 255. Alert and Block thresholds specified in the Data Filtering Profile are a function of this weight.<br>• **CC#**—Specify a weight for the credit card field (range 0-255).<br>• **SSN#**—Specify a weight for the social security number field, where the field includes dashes, such as 123-45-6789 (range 0-255, 255 is highest weight).<br>• **SSN# (without dash)**—Specify a weight for the social security number field, where the entry is made without dashes, such as 123456789 (range 0-255, 255 is highest weight). |
| Custom Patterns | The pre-defined patterns include credit card number and social security number (with and without dashes).<br>Click **Add** to add a new pattern. Specify a name for the pattern, enter the regular expression that defines the pattern, and enter a weight to assign to the pattern. Add additional patterns as needed. |

## Defining Spyware and Vulnerability Signatures

▶ *Objects > Custom Signatures > Spyware & Vulnerabilities*

The firewall supports the ability to create custom spyware and vulnerability signatures using the firewall threat engine. You can write custom regular expression patterns to identify spyware phone home communication or vulnerability exploits. The resulting spyware and vulnerability patterns become available for use in any custom vulnerability profiles. The firewall looks for the custom-defined patterns in network traffic and takes the specified action for the vulnerability exploit. Support is provided for creation of custom signatures using HTTP, SMTP, IMAP, FTP, POP3, SMB, MSSQL, MSRPC, RTSP, SSH, SSL, Telnet, Unknown-TCP, and Unknown-UDP.

You can optionally include a time attribute when defining custom signatures by specifying a threshold per interval for triggering possible actions in response to an attack. Action is taken only after the threshold is reached.

Use the **Custom Signatures** page to define signatures for vulnerability profiles.

**Table 108. Custom Signatures - Vulnerability and Spyware**

| Field | Description |
|---|---|
| **Configuration Tab** | |
| Threat ID | Enter a numeric identifier for the configuration. For spyware signatures, the range is 15000-18000; for vulnerability signatures the range is 41000-45000. |
| Name | Specify the threat name. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |
| Comment | Enter an optional comment. |

**Table 108. Custom Signatures - Vulnerability and Spyware (Continued)**

| Field | Description |
|---|---|
| Severity | Assign a level that indicates the seriousness of the threat. |
| Default Action | Assign the default action to take if the threat conditions are met:<br>• **Alert**—Generate an alert.<br>• **Drop Packets**—Do not allow packets through.<br>• **Reset Both**—Reset the client and server.<br>• **Reset Client**—Reset the client.<br>• **Reset Server**—Reset the server.<br>• **Block IP**—Block traffic for a specified period of time. Choose whether to block traffic for the source only or source and destination, and enter the duration (seconds). |
| Direction | Indicate whether the threat is assessed from the client to server, server to client, or both. |
| Affected System | Indicate whether the threat involves the client, server, either, or both. Applies to vulnerability signatures, but not spyware signatures. |
| CVE | Specify the common vulnerability enumeration (CVE) as an external reference for additional background and analysis. |
| Vendor | Specify the vendor identifier for the vulnerability as an external reference for additional background and analysis. |
| Bugtraq | Specify the bugtraq (similar to CVE) as an external reference for additional background and analysis. |
| Reference | Add any links to additional analysis or background information. The information is shown when a user clicks on the threat from the ACC, logs, or vulnerability profile. |
| **Signatures Tab** | |
| Standard Signature | Select the **Standard** radio button and then click **Add** to add a new signature. Specify the following information:<br>• **Standard**—Enter a name to identify the signature.<br>• **Comment**—Enter an optional description.<br>• **Ordered Condition Match**—Select if the order in which signature conditions are defined is important.<br>• **Scope**—Select whether to apply this signature only to the current transaction or to the full user session.<br>Specify conditions to define signatures:<br>• Add a condition by clicking **Add AND Condition** or **Add OR Condition**. To add a condition within a group, select the group and then click **Add Condition**. Select from the **Method** and **Context** drop-down lists. Specify a regular expression in the **Pattern** field. Add additional patterns as needed.<br>• To move a condition within a group, select the condition and click the **Move Up** or **Move Down** arrow. To move a group, select the group and click the **Move Up** or **Move Down** arrow. You cannot move conditions from one group to another. |

**Table 108. Custom Signatures - Vulnerability and Spyware (Continued)**

| Field | Description |
|---|---|
| Combination Signature | Select the **Combination** radio button. In the area above the subtabs, specify the following information:<br>On the **Combination Signatures** subtab, specify conditions to define signatures:<br>• Add a condition by clicking **Add AND Condition** or **Add OR Condition**. To add a condition within a group, select the group and then click **Add Condition**. Select from the **Method** and **Context** drop-down lists. Specify a regular expression in the **Pattern** field. Add additional patterns as needed.<br>• To move a condition within a group, select the condition and click the **Move Up** or **Move Down** arrow. To move a group, select the group and click the **Move Up** or **Move Down** arrow. You cannot move conditions from one group to another.<br>On the **Time Attribute** subtab, specify the following information:<br>• **Number of Hits**—Specify the threshold that will trigger any policy-based action as a number of hits (1-1000) in a specified number of seconds (1-3600).<br>• **Aggregation Criteria**—Specify whether the hits are tracked by source IP address, destination IP address, or a combination of source and destination IP addresses. |

# Security Profile Groups

▶ *Objects > Security Profile Groups*

The firewall supports the ability to create security profile groups, which specify sets of security profiles that can be treated as a unit and then added to security policies. For example, you can create a "threats" security profile group that includes profiles for antivirus, anti-spyware, and vulnerability and then create a security policy that includes the "threats" profile.

Antivirus, anti-spyware, vulnerability protection, URL filtering, and file blocking profiles that are often assigned together can be combined into profile groups to simplify the creation of security policies.

To define new security profiles, refer to "Defining Security Policies" on page 187.

**Table 109. Security Profile Group Settings**

| Field | Description |
|---|---|
| Name | Enter the profile group name (up to 31 characters). This name appears in the profiles list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems. |

**Table 109. Security Profile Group Settings (Continued)**

| Field | Description |
|---|---|
| Profiles | Select an antivirus, anti-spyware, vulnerability protection, URL filtering, and/or file blocking profile to be included in this group. Data filtering profiles can also be specified in security profile groups. Refer to "Data Filtering Profiles" on page 223. |

# Log Forwarding

▶ *Objects > Log Forwarding*

Each security policy can specify a log forwarding profile that determines whether traffic and threat log entries are logged remotely with Panorama, and/or sent as SNMP traps, syslog messages, or email notifications. By default, only local logging is performed.

Traffic logs record information about each traffic flow, and threat logs record the threats or problems with the network traffic, such as virus or spyware detection. Note that the antivirus, anti-spyware, and vulnerability protection profiles associated with each rule determine which threats are logged (locally or remotely). To apply logging profiles to security policies, refer to "Security Policies" on page 187.

**Table 110. Log Forwarding Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a profile name (up to 31 characters). This name appears in the list of log forwarding profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems. |
| **Traffic Settings** | |
| Panorama | Select the check box to enable sending traffic log entries to the Panorama centralized management system. To define the Panorama server address, refer to "Defining Management Settings" on page 30. |
| SNMP Trap<br>Email<br>Syslog | Select the SNMP, syslog, and/or email settings that specify additional destinations where the traffic log entries are sent. To define new destinations, refer to:<br>• "Configuring SNMP Trap Destinations" on page 75.<br>• "You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in Device > Server Profiles > Syslog. Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type." on page 77<br>• "Configuring Syslog Servers" on page 76 |

**Table 110. Log Forwarding Profile Settings (Continued)**

| Field | Description |
|---|---|
| **Threat Log Settings** | |
| Panorama | Click the check box for each severity level of the threat log entries to be sent to Panorama. The severity levels are:<br>• **Critical**—Very serious attacks detected by the threat security engine.<br>• **High**—Major attacks detected by the threat security engine.<br>• **Medium**—Minor attacks detected by the threat security engine, including URL blocking.<br>• **Low**—Warning-level attacks detected by the threat security engine.<br>• **Informational**—All other events not covered by the other severity levels, including informational attack object matches. |
| SNMP Trap<br>Email<br>Syslog | Under each severity level, select the SNMP, syslog, and/or email settings that specify additional destinations where the threat log entries are sent. |

# Decryption Profiles

▶ *Objects > Decryption Profiles*

Decryption profiles enable you to block and control specific aspects of the SSL forward proxy, SSL inbound inspection, and SSH traffic. After you create a decryption profile, you can then apply that profile to a decryption policy.
You can also control the trusted CAs that your device trusts, for more information, refer to "Default Trusted Certificate Authorities" on page 88.

**Table 111. Decryption Profile Settings**

| Field | Description |
|---|---|
| **SSL Forward Proxy Tab** | |
| Server Certificate Checks | Select options to control server certificates. |
| Block sessions with expired certificates | Terminate the SSL connection if the server certificate is expired. This will prevent a user from being able to accept an expired certificate and continuing with an SSL session. |
| Block sessions with untrusted issuers | Terminate the SSL session if the server certificate issuer is untrusted. |
| Restrict certificate extensions | Limits the certificate extensions used in the dynamic server certificate to key usage and extended key usage.<br>Details—Displays details on the values used for key usage and extended key usage. |
| Unsupported Mode Checks | Select options to control unsupported SSL applications. |
| Block sessions with unsupported version | Terminate sessions if the "client hello" message is not supported by PAN-OS. The SSL versions supported by PAN-OS are: SSLv3, TLS1.0, and TLS1.1. |

**Table 111. Decryption Profile Settings**

| Field | Description |
|---|---|
| Block sessions with unsupported cipher suites | Terminate the session if the cipher suite specified in the SSL handshake if it is not supported by PAN-OS. |
| Block sessions with client authentication | Terminate sessions with client authentication for SSL forward proxy traffic. |
| Failure Checks | Select the action to take if system resources are not available to process decryption. |
| Block sessions if resources not available | Terminate sessions if system resources are not available to process decryption. |
| | ***Note:*** *For unsupported modes and failure modes, the session information is cached for 12 hours, so future sessions between the same hosts and server pair are not decrypted. Use the check boxes to block those sessions instead.* |
| **SSL Inbound Inspection Tab** | |
| Unsupported Mode Checks | Selection options to control sessions if unsupported modes are detected in SSL traffic. |
| Block sessions with unsupported versions | Terminate sessions if the "client hello" message is not supported by PAN-OS. The SSL versions supported by PAN-OS are: SSLv3, TLS1.0, and TLS1.1. |
| Block sessions with unsupported cipher suites | Terminate the session if the cipher suite used is not supported by PAN-OS. |
| Failure Checks | Select the action to take if system resources are not available. |
| Block sessions if resources not available | Terminate sessions if system resources are not available to process decryption. |
| **SSH Tab** | |
| Unsupported Mode Checks | Selection options to control sessions if unsupported modes are detected in SSH traffic. Supported SSH version is SSH version 2. |
| Block sessions with unsupported versions | Terminate sessions if the "client hello" message is not supported by PAN-OS. |
| Block sessions with unsupported algorithms | Terminate sessions if the algorithm specified by the client or server is not supported by PAN-OS. |
| Failure Checks | Select actions to take if SSH application errors occur and if system resources are not available. |

**Table 111. Decryption Profile Settings**

| Field | Description |
|---|---|
| Block sessions on SSH errors | Terminate sessions if SSH errors occur. |
| Block sessions if resources not available | Terminate sessions if system resources are not available to process decryption. |

# Schedules

▶ *Objects > Schedules*

By default, each security policy applies to all dates and times. To limit a security policy to specific times, you can define schedules, and then apply them to the appropriate policies. For each schedule, you can specify a fixed date and time range or a recurring daily or weekly schedule. To apply schedules to security policies, refer to "Security Policies" on page 187.

***Note:*** *When a security policy is invoked by a defined schedule, only new sessions are affected by the applied security policy. Existing sessions are not affected by the scheduled policy.*

**Table 112. Schedule Settings**

| Field | Description |
|---|---|
| Name | Enter a schedule name (up to 31 characters). This name appears in the schedule list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems. |
| Recurrence | Select the type of schedule (**Daily**, **Weekly**, or **Non-Recurring**). |
| Daily | Click **Add** and specify a start and end time in 24-hour format (HH:MM). |
| Weekly | Click **Add**, select a day of the week, and specify the start and end time in 24-hour format (HH:MM). |
| Non-recurring | Click **Add** and specify a start and end date and time. |

## Chapter 6
# Reports and Logs

This chapter describes how to view the reports and logs provided with the firewall:


- "Using the Dashboard" in the next section
- "Using the Application Command Center" on page 253
- "Using App-Scope" on page 256
- "Viewing the Logs" on page 264
- "Working with Botnet Reports" on page 267
- "Managing PDF Summary Reports" on page 270
- "Managing User Activity Reports" on page 272
- "Managing Report Groups" on page 272
- "Scheduling Reports for Email Delivery" on page 273
- "Viewing Reports" on page 273
- "Generating Custom Reports" on page 274
- "Identifying Unknown Applications and Taking Action" on page 276
- "Taking Packet Captures" on page 278

***Note:*** *Most of the reports in this section support optional selection of a virtual system from the drop-down list at the top of page.*

# Using the Dashboard

▶ *Dashboard*

The **Dashboard** page displays general device information, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available charts are displayed by default, but each user can remove and add individual charts, as needed.

Click **Refresh** to update the Dashboard. To change the automatic refresh interval, select an interval from the drop-down list (1 min, 2 mins, 5 mins, or Manual). To add a Widget to the Dashboard, click the Widget drop-down, select a category and then the widget name. To delete a widget, click in the title bar.

Review the following information in each chart.

**Table 113. Dashboard Charts**

| Chart | Description |
|---|---|
| Top Applications | Displays the applications with the most sessions. The block size indicates the relative number of sessions (mouse-over the block to view the number), and the color indicates the security risk—from green (lowest) to red (highest). Click an application to view its application profile. |
| Top High Risk Applications | Similar to Top Applications, except that it displays the highest-risk applications with the most sessions. |
| General Information | Displays the device name, model, PAN-OS software version, the application, threat, and URL filtering definition versions, the current date and time, and the length of time since the last restart. |
| Interface Status | Indicates whether each interface is up (green), down (red), or in an unknown state (gray). |
| Threat Logs | Displays the threat ID, application, and date and time for the last 10 entries in the Threat log. The threat ID is a malware description or URL that violates the URL filtering profile. |
| Config Logs | Displays the administrator user name, client (Web or CLI), and date and time for the last 10 entries in the Configuration log. |
| Data Filtering Logs | Displays the description and date and time for the last 60 minutes in the Data Filtering log. |
| URL Filtering Logs | Displays the description and date and time for the last 60 minutes in the URL Filtering log. |
| System Logs | Displays the description and date and time for the last 10 entries in the System log. Note that a "Config installed" entry indicates configuration changes were committed successfully. |
| System Resources | Displays the Management CPU usage, Data Plane usage, and the Session Count, which displays the number of sessions established through the firewall. |
| Logged In Admins | Displays the source IP address, session type (Web or CLI), and session start time for each administrator who is currently logged in. |
| ACC Risk Factor | Displays the average risk factor (1 to 5) for the network traffic processed over the past week. Higher values indicate higher risk. |

**Table 113. Dashboard Charts (Continued)**

| Chart | Description |
|---|---|
| High Availability | If high availability (HA) is enabled, indicates the HA status of the local and peer device—green (active), yellow (passive), or black (other). For more information about HA, refer to "Enabling HA on the Firewall" on page 101. |
| Locks | Shows configuration locks taken by administrators. |

# Using the Application Command Center

▶ *ACC*

The **Application Command Center (ACC)** page displays the overall risk level for your network traffic, the risk levels and number of threats detected for the most active and highest-risk applications on your network, and the number of threats detected from the busiest application categories and from all applications at each risk level. The ACC can be viewed for the past hour, day, week, month, or any custom-defined time frame.

Risk levels (1=lowest to 5=highest) indicate the application's relative security risk based on criteria such as whether the application can share files, is prone to misuse, or tries to evade firewalls.

To view the Application Command Center:

1. Under the **ACC** tab, change one or more of the following settings at the top of the page, and click **Go**:
   a. Select a virtual system, if virtual systems are defined.
   b. Select a time period from the **Time** drop-down list. The default is Last Hour.
   c. Select a sorting method from the **Sort By** drop-down list. You can sort the charts in descending order by number of sessions, bytes, or threats. The default is by number of sessions.
   d. For the selected sorting method, select the top number of applications and application categories shown in each chart from the **Top** drop-down list.

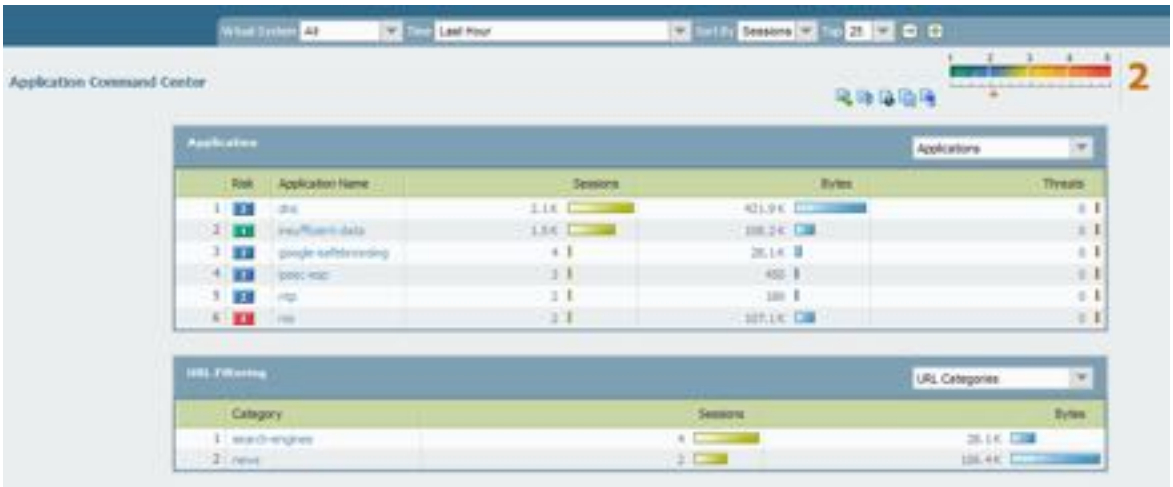

**Figure 22. Application Command Center Page**

2. To open log pages associated with the information on the page, use the log links in the upper-right corner of the page, as shown here. The context for the logs matches the information on the page.

3. To filter the list, click an item in one of the columns, this will add that item to the filter bar located above the log column names. After adding the desired filters, click the Apply Filter icon.

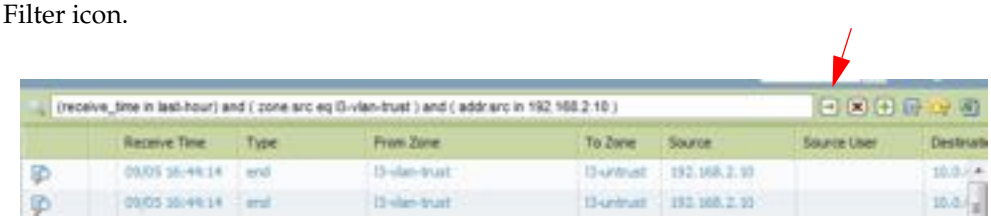


4. Choose a view from the drop-down list for the area of interest, as described in the following table.

5. Use the drop-down lists for Applications, URL Categories, Threats, Content/File Types, and HIP Objects.

**Table 114. Application Command Center Charts**

| Chart | Description |
|---|---|
| Application | Displays information organized according to the menu selection. Information includes the number of sessions, bytes transmitted and received, number of threats, application category, application subcategories, application technology, and risk level, as applicable.<br>• Applications<br>• High risk applications<br>• Categories<br>• Sub Categories<br>• Technology<br>• Risk |
| URL Filtering | Displays information organized according to the menu selection. Information includes the URL, URL category, repeat count (number of times access was attempted, as applicable).<br>• URL Categories<br>• URLs<br>• Blocked URL Categories<br>• Blocked URLs |

**Table 114. Application Command Center Charts (Continued)**

| Chart | Description |
|---|---|
| Threat Prevention | Displays information organized according to the menu selection. Information includes threat ID, count (number of occurrences), number of sessions, and subtype (such as vulnerability), as applicable.<br>• Threats<br>• Types<br>• Spyware<br>• Spyware Phone Home<br>• Spyware Downloads<br>• Vulnerability<br>• Virus |
| Data Filtering | • Content/File Types<br>• Types<br>• File Names |
| HIP Matches | • HIP Objects<br>• HIP Profiles |

6. To view additional details, click any of the links. A details page opens to show information about the item at the top and additional lists for related items.

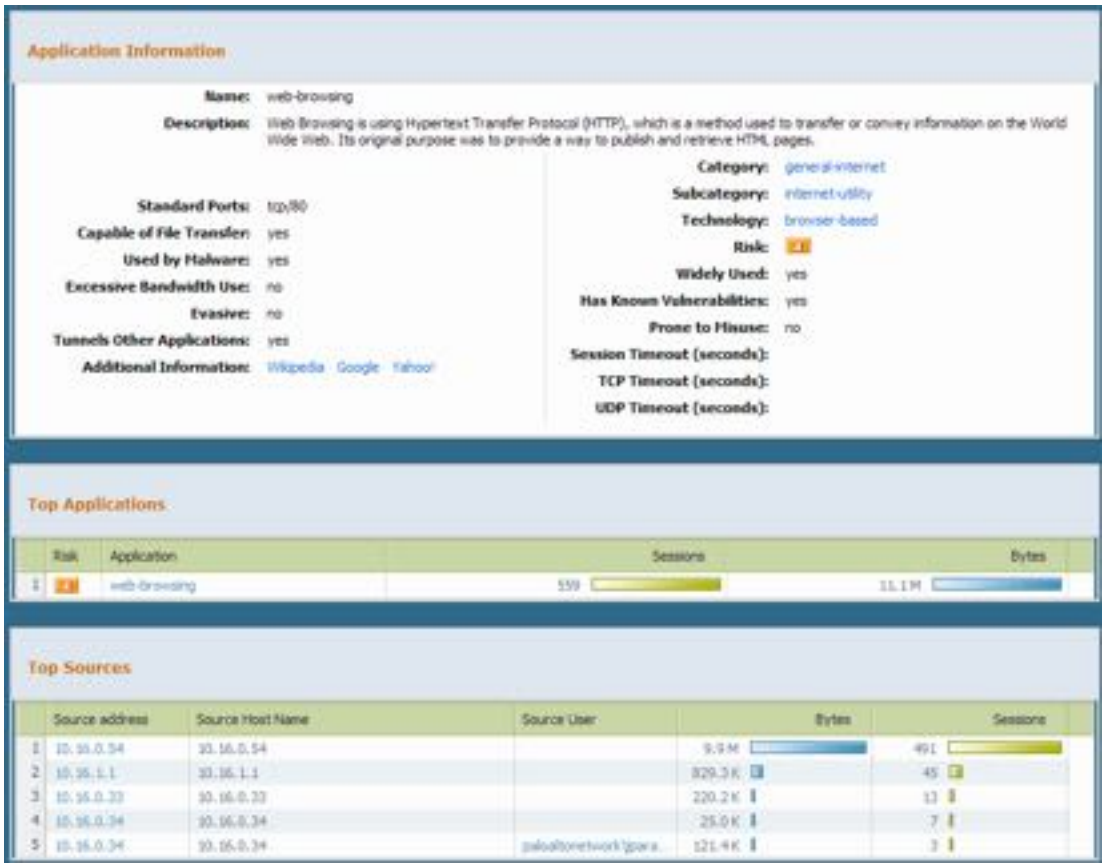


**Figure 23. Application Command Center Drill Down Page**

# Using App-Scope

▶ *Monitor > App Scope*

The App-Scope reports introduce a visibility and analysis tools to help pinpoint problematic behavior, helping you understand the following aspects of your network:

- Changes in application usage and user activity
- Users and applications that take up most of the network bandwidth
- Network threats

With the App-Scope reports, you can quickly see if any behavior is unusual or unexpected. Each report provides a dynamic, user-customizable window into the network. The reports include options to select the data and ranges to display.

To view the reports, click the report name under **App-Scope** on the left side of the page in the **Monitor** tab. Select one of the report types lists below. Report options are available from the drop-down lists at the top and bottom of some of the pages.

**Table 115. Application Command Center Charts**

| Chart | Description |
|---|---|
| Summary | "Summary Report" on page 257 |
| Change Monitor | "Change Monitor Report" on page 258 |
| Threat Monitor | "Threat Monitor Report" on page 259 |
| Threat Map | "Threat Monitor Report" on page 259 |
| Network Monitor | "Network Monitor Report" on page 261 |
| Traffic Map | "Traffic Map Report" on page 263 |

## Summary Report

The Summary report (Figure 24) displays charts for the top five gainers, losers, and bandwidth consuming applications, application categories, users, and sources.

**Figure 24. App-Scope Summary Report**

## Change Monitor Report

The Change Monitor report (Figure 25) displays changes over a specified time period. For example, Figure 25 displays the top applications that gained in use over the last hour as compared with the last 24-hour period. The top applications are determined by session count and sorted by per cent.

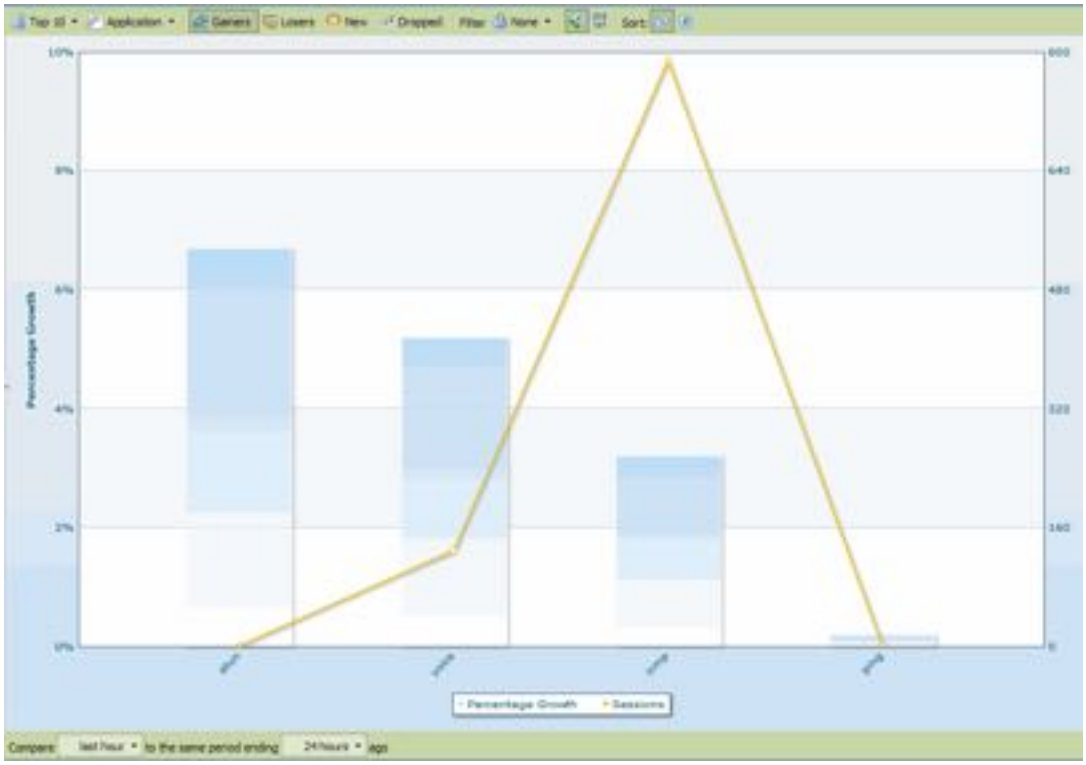

**Figure 25. App-Scope Change Monitor Report**

This report contains the following buttons and options.

**Table 116. Change Monitor Report Options**

| Item | Description |
|---|---|
| **Top Bar** | |
| Top 10 | Determines the number of records with the highest measurement included in the chart. |
| Application | Determines the type of item reported: Application, Application Category, Source, or Destination. |
| Gainers | Displays measurements of items that have increased over the measured period. |
| Losers | Displays measurements of items that have decreased over the measured period. |
| New | Displays measurements of items that were added over the measure period. |

**Table 116. Change Monitor Report Options (Continued)**

| Item | Description |
|---|---|
| zzZ Dropped | Displays measurements of items that were discontinued over the measure period. |
| Filter None | Applies a filter to display only the selected item. **None** displays all entries. |
| 010 101 | Determines whether to display session or byte information. |
| Sort: % # | Determines whether to sort entries by percentage or raw growth. |
| **Bottom Bar** | |
| Compare last hour to the same period ending 24 hours ago | Specifies the period over which the change measurements are taken. |

## Threat Monitor Report

The Threat Monitor report (Figure 26) displays a count of the top threats over the selected time period. For example, Figure 26 shows the top 10 threat types for the past 6 hours.

**Figure 26. App-Scope Threat Monitor Report**

Each threat type is color-coded as indicated in the legend below the chart. This report contains the following buttons and options.

**Table 117. Threat Monitor Report Buttons**

| Button | Description |
|---|---|
| **Top Bar** | |
| | Determines the number of records with the highest measurement included in the chart. |
| | Determines the type of item measured: Threat, Threat Category, Source, or Destination. |
| | Applies a filter to display only the selected type of items. |
| | Determines whether the information is presented in a stacked column chart or a stacked area chart. |
| **Bottom Bar** | |
| | Specifies the period over which the measurements are taken. |

## Threat Map Report

The Threat Map report (Figure 27) shows a geographical view of threats, including severity.

**Figure 27. App-Scope Threat Monitor Report**

Each threat type is color-coded as indicated in the legend below the chart. Click a country on the map to zoom in. Click the **Zoom Out** button in the lower right corner of the screen to zoom out. This report contains the following buttons and options.

**Table 118. Threat Map Report Buttons**

| Button | Description |
|---|---|
| **Top Bar** | |
| | Determines the number of records with the highest measurement included in the chart. |
| 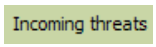 | Displays incoming threats. |
| 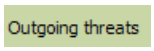 | Displays outgoing threats. |
| | Applies a filter to display only the selected type of items. |
| **Bottom Bar** | |
| 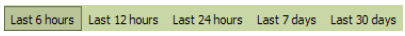 | Indicates the period over which the measurements are taken. |

## Network Monitor Report

The Network Monitor report (Figure 28) displays the bandwidth dedicated to different network functions over the specified period of time. Each network function is color-coded as indicated in the legend below the chart. For example, Figure 28 shows application bandwidth for the past 7 days based on session information.

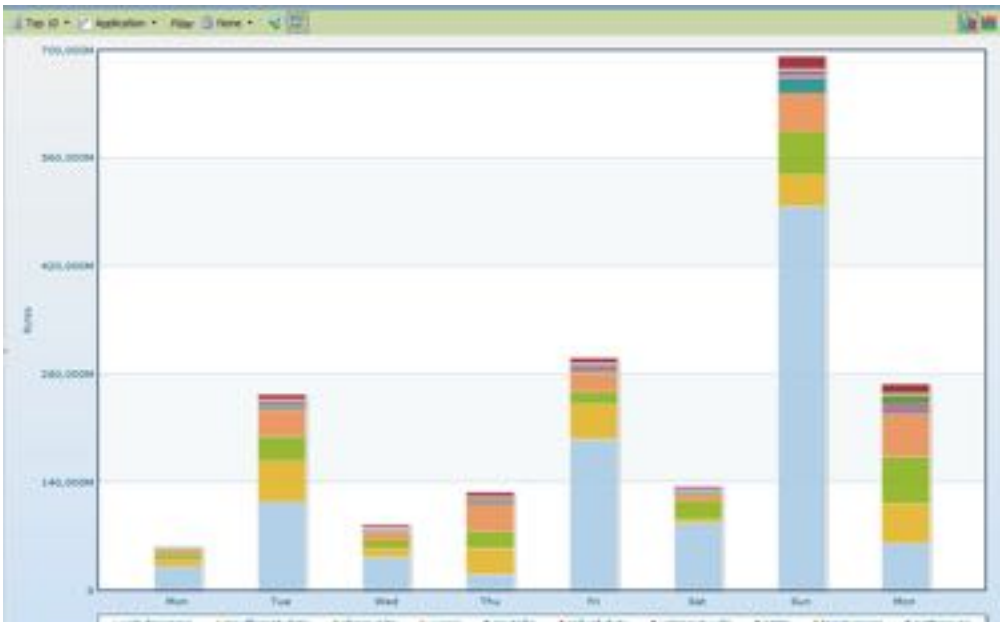

**Figure 28. App-Scope Network Monitor Report**

The report contains the following buttons and options.

**Table 119. Network Monitor Report Buttons**

| Button | Description |
|---|---|
| **Top Bar** | |
| | Determines the number of records with the highest measurement included in the chart. |
| 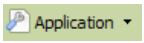  | Determines the type of item reported: Application, Application Category, Source, or Destination. |
| 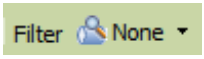  | Applies a filter to display only the selected item. **None** displays all entries. |
| | Determines whether to display session or byte information. |
| | Determines whether the information is presented in a stacked column chart or a stacked area chart. |
| **Bottom Bar** | |
| | Indicates the period over which the change measurements are taken. |

## Traffic Map Report

The Traffic Map report (Figure 29) shows a geographical view of traffic flows according to sessions or flows.

**Figure 29. App-Scope Traffic Monitor Report**

Each traffic type is color-coded as indicated in the legend below the chart. This report contains the following buttons and options.

**Table 120. Threat Map Report Buttons**

| Button | Description |
|---|---|
| **Top Bar** | |
| Top 10 | Determines the number of records with the highest measurement included in the chart. |
| Incoming threats | Displays incoming threats. |
| Outgoing threats | Displays outgoing threats. |
| 010 101 | Determines whether to display session or byte information. |

**Table 120. Threat Map Report Buttons (Continued)**

| Button | Description |
|---|---|
| **Bottom Bar** | |
| Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days | Indicates the period over which the change measurements are taken. |

# Viewing the Logs

▶ *Monitor > Logs*

The firewall maintains logs for WildFire, configurations, system, alarms, traffic flows, threats, URL filtering, data filtering, and Host Information Profile (HIP) matches. You can view the current logs at any time. To locate specific entries, you can apply filters to most of the log fields.

***Note:*** *The firewall displays the information in logs so that role-based administration permissions are respected. When you display logs, only the information that you have permission to see is included. For information on administrator permissions, refer to "Defining Administrator Roles" on page 58.*

To view the logs, click the log types on the left side of the page in the **Monitor** tab.

Each log page has a filter area at the top of the page.

Use the filter area as follows:

- Click any of the underlined links in the log listing to add that item as a log filter option. For example, if you click the **Host** link in the log entry for 10.0.0.252 and **Web Browsing** in both items are added, and the search will find entries that match both (AND search).

***Note:*** *To filter information for a user group, add the query (user in 'domain/user-group'). For example, to generate a list of the activities for the marketing group in Company ABC, you would query (user in 'company_abc/marketing').*

- To define other search criteria, click the **Add Log Filter** icon. Select the type of search (and/or), the attribute to include in the search, the matching operator, and the values for the match, if appropriate. Click **Add** to add the criterion to the filter area on the log page, and then click **Close** to close the pop-up window. Click the **Apply Filter** icon to display the filtered list.

***Note:*** *You can combine filter expressions added on the log page with those that you define in the Expression pop-up window. Each is added as an entry on the Filter line on the log page.*

*If you set the "in" Received Time filter to* ***Last 60 seconds****, some of the page links on the log viewer may not show results because the number of pages may grow or shrink due to the dynamic nature of the selected time.*

- To clear filters and redisplay the unfiltered list, click the **Clear Filter** button.
- To save your selections as a new filter, click the **Save Filter** button, enter a name for the filter, and click **OK**.
- To export the current log listing (as shown on the page, including any applied filters) click the **Save Filter** button. Select whether to open the file or save it to disk, and select the check box if you want to always use the same option. Click **OK**.

To change the automatic refresh interval, select an interval from the drop-down list (1 min, 30 seconds, 10 seconds, or Manual). To change the number of log entries per page, select the number of rows from the **Rows** drop-down list.

Log entries are retrieved in blocks of 10 pages. Use the paging controls at the bottom of the page to navigate through the log list. Select the **Resolve Hostname** check box to begin resolving external IP addresses to domain names.

To display additional details, click the spyglass icon for an entry.

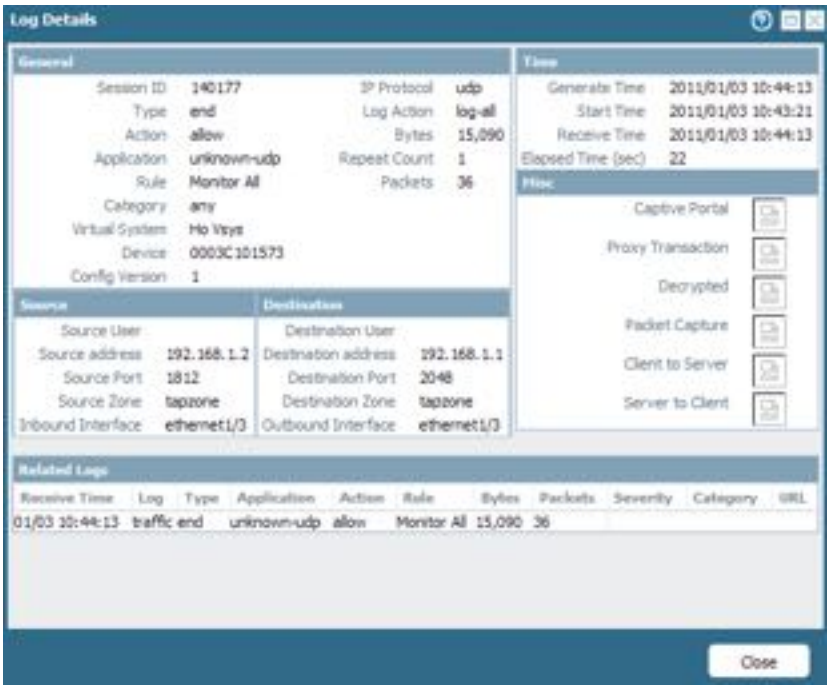


**Figure 30. Log Entry Details**

If the source or destination has an IP address to name mapping defined in the **Addresses** page, the name is presented instead of the IP address. To view the associated IP address, move your cursor over the name.

Review the following information in each log.

**Table 121. Log Descriptions**

| Chart | Description |
|---|---|
| Traffic | Displays an entry for the start and end of each session. Each entry includes the date and time, the source and destination zones, addresses, and ports, the application name, the security rule name applied to the flow, the rule action (allow, deny, or drop), the ingress and egress interface, and the number of bytes.<br>Click next to an entry to view additional details about the session, such as whether an ICMP entry aggregates multiple sessions between the same source and destination (the Count value will be greater than one).<br>Note that the **Type** column indicates whether the entry is for the start or end of the session, or whether the session was denied or dropped. A "drop" indicates that the security rule that blocked the traffic specified "any" application, while a "deny" indicates the rule identified a specific application.<br>If traffic is dropped before the application is identified, such as when a rule drops all traffic for a specific service, the application is shown as "not-applicable". |
| Threat | Displays an entry for each security alarm generated by the firewall. Each entry includes the date and time, a threat name or URL, the source and destination zones, addresses, and ports, the application name, and the alarm action (allow or block) and severity.<br>Click next to an entry to view additional details about the threat, such as whether the entry aggregates multiple threats of the same type between the same source and destination (the Count value will be greater than one).<br>Note that the **Type** column indicates the type of threat, such as "virus" or "spyware." The **Name** column is the threat description or URL, and the **Category** column is the threat category (such as "keylogger") or URL category.<br>If local packet captures are enabled, click next to an entry to access the captured packets, as in the following figure. To enable local packet captures, refer to the subsections under "Security Profiles" on page 211. |
| URL Filtering | Displays logs for URL filters, which block access to specific web sites and web site categories or generate an alert when a proscribed web site is accessed. Refer to "URL Filtering Profiles" on page 217 for information on defining URL filtering profiles. |
| WildFire | Displays logs for files that are uploaded and analyzed by the WildFire server, log data is sent back to the device after analysis, along with the analysis results.<br>A subscription is required for this feature. If you do not have a subscription, you can use the WildFire Portal to view log information. Refer to "Using the WildFire Portal" on page 425. |
| Data Filtering | Displays logs for the security policies that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall. Refer to "Data Filtering Profiles" on page 223 for information on defining data filtering profiles.<br>To configure password protection for access the details for a log entry, click the icon. Enter the password and click **OK**. Refer to "Defining Custom Response Pages" on page 115 for instructions on changing or deleting the data protection password.<br>***Note:*** *The system prompts you to enter the password only once per session.* |

**Table 121. Log Descriptions (Continued)**

| Chart | Description |
|---|---|
| Configuration | Displays an entry for each configuration change. Each entry includes the date and time, the administrator user name, the IP address from where the change was made, the type of client (Web or CLI), the type of command executed, whether the command succeeded or failed, the configuration path, and the values before and after the change. |
| System | Displays an entry for each system event. Each entry includes the date and time, the event severity, and an event description. |
| HIP Match | Displays information about security policies that apply to GlobalProtect clients. For more information, refer to "Overview" on page 335. |
| Alarms | The alarms log records detailed information on alarms that are generated by the system. The information in this log is also reported in the **Alarms** window. Refer to "Viewing Alarms" on page 85. |

### Viewing Session Information

▶ *Monitor > Session Browser*

Open the **Session Browser** page to browse and filter current running sessions on the firewall. For information on filtering options for this page, refer to "Viewing the Logs" on page 264.

## Working with Botnet Reports

The botnet report feature allows you to use behavior-based mechanisms to identify potential botnet-infected hosts in the network. Using network, threat, URL, and data filtering logs, the firewall evaluates threats based on criteria that include visits to malware sites and dynamic DNS sites, visits to recently registered domains (within the last 30 days), unknown application usage, and the existence of Internet Relay Chat (IRC) traffic.

After correlating and identifying hosts that match infected botnet behavior, the firewall assigns each potentially infected host a confidence score of 1 to 5 to indicate the likelihood of botnet infection (1 indicates the lowest and 5 the highest likelihood of infection). Because behavior-based detection mechanisms require correlating traffic across multiple logs over a period of 24 hours, the firewall generates a report every 24 hours that contains a sorted list of hosts based on confidence level.

## Configuring the Botnet Report

▶ *Monitor > Botnet*

Use these settings to specify types of suspicious traffic (traffic that may indicate botnet activity). To configure the settings, click the **Configuration** button on the right side of the **Botnet** page.

**Table 122. Botnet Configuration Settings**

| Field | Description |
|---|---|
| HTTP Traffic | Select the **Enable** check box for the events that you want to include in the reports:<br>• **Malware URL visit**—Identifies users communicating with known malware URLs based on malware and botnet URL filtering categories.<br>• **Use of dynamic DNS**—Looks for dynamic DNS query traffic that could indicate botnet communication.<br>• **Browsing to IP domains**—Identifies users that browse to IP domains instead of URLs.<br>• **Browsing to recently registered domains**—Looks for traffic to domains that have been registered within the past 30 days.<br>• **Executable files from unknown sites**—Identifies executable files downloaded from unknown URLs. |
| Unknown Applications | Select the check boxes to mark unknown TCP or unknown UDP applications as suspicious, and specify the following information:<br>• **Sessions Per Hour**—Number of application sessions per hour that are associated with the unknown application.<br>• **Destinations Per Hour**—Number of destinations per hour that are associated with the unknown application.<br>• **Minimum Bytes**—Minimum payload size<br>• **Maximum Bytes**—Maximum payload size. |
| IRC | Select the check box to include IRC servers as suspicious. |

## Managing Botnet Reports

▶ *Monitor > Botnet > Report Setting*

You can specify report queries and then generate and view botnet analysis reports. The reports are based on botnet configuration settings (refer to "Configuring the Botnet Report" on page 268). You can include or exclude source or destination IP addresses, users, zones, interfaces, regions, or countries.

Scheduled reports run once per day. You can also generate and display reports on demand by clicking **Run Now** in the window where you define the report queries. The generated report is displayed on the **Botnet** page.

To manage botnet reports, click the **Report Setting** button on the right side of the **Botnet** page.

To export a report, select the report and click **Export to PDF** or **Export to CSV**.

**Table 123. Botnet Report Settings**

| Field | Description |
|---|---|
| Test Run Time Frame | Select the time interval for the report (last 24 hours or last calendar day). |
| # Rows | Specify the number of rows in the report. |
| Scheduled | Select the check box to run the report on a daily basis. To run the report manually, click **Run Now** at the top of the Botnet Report window. |
| Query | Construct the report query by specifying the following, and then clicking **Add** to add the configured expression to the query. Repeat as needed to construct the complete query:<br>• **Connector**—Specify a logical connector (AND/OR).<br>• **Attribute**—Specify the source or destination zone, address, or user.<br>• **Operator**—Specify the operator to relate the attribute to a value.<br>• **Value**—Specify the value to match. |
| Negate | Select the check box to apply the negation of the specified query, meaning that the report will contain all information that is not a result of the defined query. |

# Managing PDF Summary Reports

▶ *Monitor > PDF Reports > Manage PDF Summary*

PDF summary reports contain information compiled from existing reports, based on data for the top 5 in each category (instead of top 50). They also contain trend charts that are not available in other reports.

**Figure 31. PDF Summary Report**

To create PDF summary reports, click **Add**. The **Manage PDF Summary Reports** page opens to show all of the available report elements.

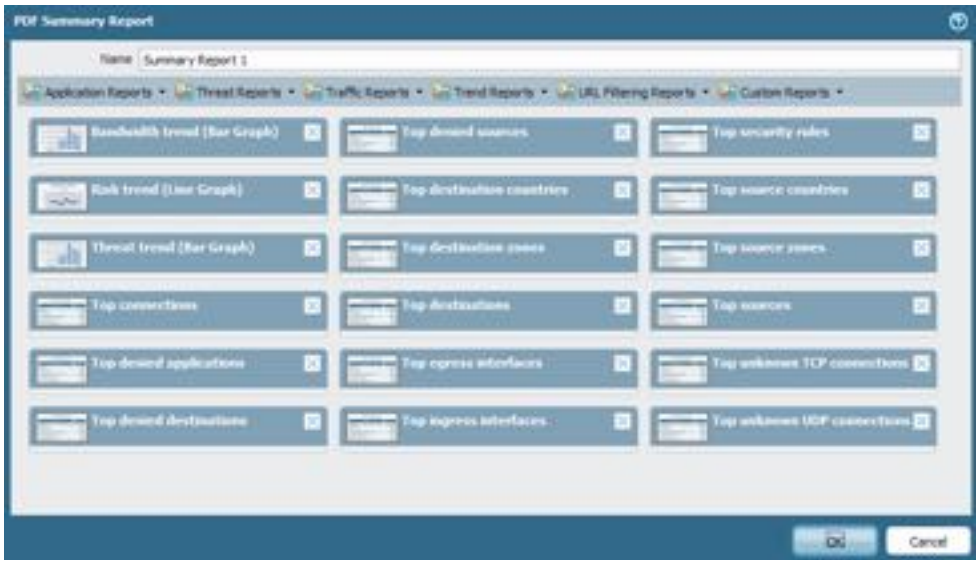


**Figure 32. Managing PDF Reports**

Use one or more of these options to design the report:

- To remove an element from the report, click the icon in the upper-right corner of the element's icon box or remove the check box from the item in the appropriate drop-down list box near the top of the page.
- Select additional elements by choosing from the drop-down list boxes near the top of the page.
- Drag and drop an element's icon box to move it to another area of the report.

***Note:*** *A maximum of 18 report elements is permitted. You may need to delete existing elements to add additional ones.*

Click **Save**, enter a name for the report, as prompted, and click **OK**.

To display PDF reports, choose **PDF Summary Report**, and select a report type from the drop-down list at the bottom of the page to display the generated reports of that type. Click an underlined report link to open or save the report.

# Managing User Activity Reports

▶ *Monitor > PDF Reports > User Activity Report*

Use this page to create reports that summarize the activity of individual users. Click **New** and specify the following information.

**Table 124. User Activity Report Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the report (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| User | Enter the user name or IP address (IPv4 or IPv6) of the user who will be the subject of the report. |
| Time frame | Select the time frame for the report from the drop-down list. |

To run the report on demand, select the report and click **Edit**, and then click **Run**.

# Managing Report Groups

▶ *Monitor > PDF Reports > Report Groups*

Report groups allow you to create sets of reports that the system can compile and send as a single aggregate PDF report with an optional title page and all the constituent reports included.

**Table 125. Report Group Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the report group (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Title Page | Select the check box to include a title page in the report. |
| Title | Enter the name that will appear as the report title. |
| Report selection | Select reports from the left column and click **Add** to move each report to the report group on the right. You can select Predefined, Custom, PDF Summary, and Log View report types.<br>The **Log View** report is a report type that is automatically created each time you create a custom report and uses the same name as the custom report. This report will show the logs that were used to build the contents of the custom report.<br>To include the log view data, when creating a report group, you add your custom report under the **Custom Reports** list and then add the log view report by selecting the matching report name from the **Log View** list. When you receive the report, you will see your custom report data followed by the log data that was used to create the custom report. |

To use the report group, refer to "Scheduling Reports for Email Delivery" in the next section.

# Scheduling Reports for Email Delivery

▶ *Monitor > PDF Reports > Email Scheduler*

Use the Email scheduler to schedule reports for delivery by email. Before adding a schedule, you must define report groups and an email profile. Refer to "Managing Report Groups" on page 272 and "Configuring Email Notification Settings" on page 83.

Scheduled reports begin running at 2:00 AM, and email forwarding occurs after all scheduled reports have finished running.

**Table 126. Email Scheduler Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the schedule (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Report Group | Select the report group (refer to "Managing Report Groups" on page 272). |
| Recurrence | Select the frequency at which to generate and send the report. |
| Email Profile | Select the profile that defines the email settings. Refer to "You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in Device > Server Profiles > Syslog. Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type." on page 77 for information on defining email profiles. |
| Override Recipient email(s) | Enter an optional email address to use instead of the recipient specified in the email profile. |

# Viewing Reports

▶ *Monitor > Reports*

The firewall provides various "top 50" reports of the traffic statistics for the previous day or a selected day in the previous week.

To view the reports, click the report names on the right side of the page (Custom Reports, Application Reports, Traffic Reports, Threat Reports, URL Filtering Reports, and PDF Summary Reports).

By default, all reports are displayed for the previous calendar day. To view reports for any of the previous days, select a report generation date from the **Select** drop-down list at the bottom of the page.

The reports are listed in sections. You can view the information in each report for the selected time period. To export the log in CSV format, click **Export to CSV**. To open the log information in PDF format, click **Export to PDF**. The PDF file opens in a new window. Click the icons at the top of the window to print or save the file.

# Generating Custom Reports

▶ *Monitor > Manage Custom Reports*

You can create custom reports that are optionally based on existing report templates. The reports can be run on demand or scheduled to run each night. To view previously defined reports, choose **Reports** on the side menu.

Click **Add** to create a new custom report. To base a report on an existing template, click **Load Template** and choose the template.

Specify the following settings to define the report.

**Table 127. Custom Report Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the report (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Database | Choose the database to use as the data source for the report. |
| Time Frame | Choose a fixed time frame or choose **Custom** and specify a date and time range. |
| Sort By | Choose sorting options to organize the report, including the amount of information to include in the report. The available options depend on the choice of database. |
| Group By | Choose grouping options to organize the report, including the amount of information to include in the report. The available options depend on the choice of database. |
| Scheduled | Select the check box to run the report each night. The report then becomes available by choosing **Reports** on the side menu. |
| Columns | Choose columns to include in the report from the **Available** column and use the right-facing arrows to move them to the **Selected** column. Use the up and down arrows to reorder the selected columns, and use the left-facing arrows to remove previously selected columns. |

**Table 127. Custom Report Settings (Continued)**

| Field | Description |
|---|---|
| Query and Query Builder | To build a report query, specify the following and click **Add**. Repeat as needed to construct the full query.<br>• **Connector**—Choose the connector (and/or) to precede the expression you are adding.<br>• **Attribute**—Choose a data element. The available options depend on the choice of database.<br>• **Operator**—Choose the criterion to determine whether the attribute applies (such as =). The available options depend on the choice of database.<br>• **Value**—Specify the attribute value to match.<br>For example, the following figure (based on the Traffic Log database) shows a query that matches if the traffic log entry was received in the past 24 hours and is from the "untrust" zone.<br> |
| Negate | Select the check box to interpret the query as a negation. In the previous example, the negate option causes a match on entries that are not in the past 24 hours or are not from the "untrust" zone. |

# Identifying Unknown Applications and Taking Action

There are several ways to view unknown applications using the web interface of the Palo Alto Networks devices:

- **Application Command Center (ACC)**—Unknown applications are sorted along with other applications in the ACC. Click a link for an unknown application to view the details of the application, including top sources and destinations. For top sources, click the link to look up the owner of the address.

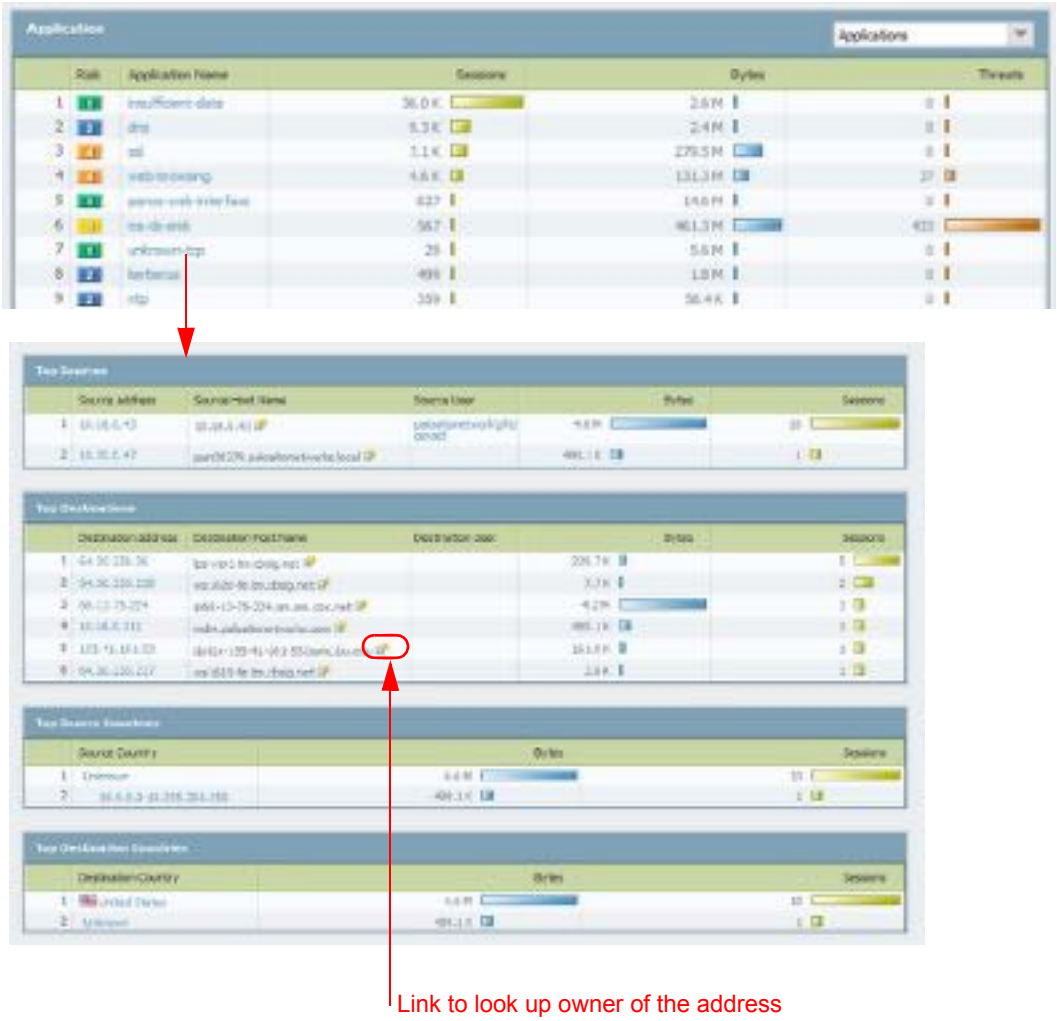


**Figure 33. Unknown Applications in the ACC List**

- **Unknown application reports**—Unknown application reports are automatically run on a daily basis and stored in the Reports section of the **Monitor** tab. These reports can provide useful information to help identify unknown applications.

- **Detailed traffic logs**—You can use the detailed traffic logs to track down unknown applications. If logging is enabled for the start and end of session, the traffic log will provide specific information about the start and end of an unknown session. Use the filter option to restrict the display to entries that match "unknown-tcp," as shown in the next figure.

**Figure 34. Unknown Applications in Traffic Logs**

## Taking Action

You can take the following actions to deal with unknown applications:

- Use custom application definition with application override (refer to "Custom Application Definition with Application Override" on page 205).
- Use custom applications with signatures (refer to "Custom Applications with Signatures" on page 236).
- Request an App-ID from Palo Alto Networks (refer to "Requesting an App-ID from Palo Alto Networks" in the next section).

Policies can also be set to control unknown applications by unknown TCP, unknown UDP or by a combination of source zone, destination zone, and IP addresses. Refer to "Application Override Policies" on page 205.

***Note:*** *You can use custom signatures in App-ID definitions.*

## Requesting an App-ID from Palo Alto Networks

If it is necessary to identify an application using application contents instead of port, protocol, and IP address, you can submit the application to Palo Alto Networks for classification. This is important for applications that run over the Internet and for which custom application does not work. You can submit the application to Palo Alto Networks in either of the following ways:

- If the application is a readily accessible on the Internet (for example, an instant messaging application), then submit the name of the application and the URL to your account team or to this web site: *http://www.paloaltonetworks.com/researchcenter/tools*.

- If the application is not easily accessible (for example, a customer relationship management application) you must submit a packet capture (PCAP) of the running application using the session packet capture function built into the firewall. For assistance, contact technical support at *support@paloaltonetworks.com*.

### Other Unknown Traffic

The firewall may report an application to be "unknown" in the ACC, logs, or reports for either of the following reasons:

- **Incomplete**—A handshake took place, but no data packets were sent prior to the timeout.
- **Insufficient-Data**—A handshake took place followed by one or more data packets; however, not enough data packets were exchanged to identify the application.

## Taking Packet Captures

▶ ***Monitor > Packet Capture***

PAN-OS supports packet capture for troubleshooting or detecting unknown applications. You can define filters such that only the packets that match the filters are captured. The packet captures are locally stored on the device and are available for download to your local computer.

***Note:*** *Packet Capture is for troubleshooting only. This feature can cause the system performance to degrade and should be used only when necessary. After the capture is complete, please remember to disable the feature.*

To specify filtering and capture options, specify the information in the following table.

To clear all filtering and capture settings, click **Clear All Settings.**

To select capture files for download, click the file name in the capture file list on the right side of the page.

**Table 128. Packet Capture Settings**

| Field | Description |
|---|---|
| **Filtering** | |
| Manage Filters | Click **Manage Filters**, click **Add** to add a new filter, and specify the following information:<br>• **Id**—Enter or select an identifier for the filter.<br>• **Ingress Interface**—Select the firewall interface.<br>• **Source**—Specify the source IP address.<br>• **Destination**—Specify the destination IP address.<br>• **Src Port**—Specify the source port.<br>• **Dest Port**—Specify the destination port.<br>• **Proto**—Specify the protocol to filter.<br>• **Non-IP**—Choose how to treat non-IP traffic (exclude all IP traffic, include all IP traffic, include only IP traffic, or do not include an IP filter).<br>• **IPv6**—Select the check box to include IPv6 packets in the filter. |

**Table 128. Packet Capture Settings (Continued)**

| Field | Description |
|---|---|
| Filtering | Click to toggle the filtering selections on or off. |
| Pre-Parse Match | Click to toggle the pre-parse match option on or off.<br>The pre-parse-match option is added for advanced troubleshooting purposes. After a packet enters the ingress port, it proceeds through several processing steps before it is parsed for matches against pre-configured filters.<br>It is possible for a packet, due to a failure, to not reach the filtering stage. This can occur, for example, if a route lookup fails.<br>Set the pre-parse-match setting to **ON** to emulate a positive match for every packet entering the system. This allows the firewall to capture even the packets that do not reach the filtering process. If a packet is able to reach the filtering stage, it is then processed according to the filter configuration and discarded if it fails to meet filtering criteria. |
| **Capture Files** | |
| Capturing | Click to toggle packet capturing on or off. |
| Capture Settings | Click **Add** and specify the following:<br>• **Stage**—Indicate the point at which to capture the packet:<br>– **drop**—When packet processing encounters an error and the packet is to be dropped.<br>– **firewall**—When the packet has a session match or a first packet with a session is successfully created.<br>– **receive**—When the packet is received on the dataplane processor.<br>– **transmit**—When the packet is to be transmitted on the dataplane processor.<br>• **File**—Specify the capture file name. The file name should begin with a letter and can include letters, digits, periods, underscores, or hyphens.<br>• **Packet Count**—Specify the number of packets after which capturing stops.<br>• **Byte Count**—Specify the number of bytes after which capturing stops. |

Chapter 7

# Configuring the Firewall for User Identification

This chapter describes how to configure the firewall to identify the users who attempt to access the network.


- "Overview of User Identification" in the next section
- "User Identification Agents" on page 284
- "PAN-OS User Mapping Configuration" on page 291
- "Setting Up the User-ID Agent" on page 296
- "Setting Up the Terminal Services Agent" on page 301


## Overview of User Identification

User Identification (User-ID) is a feature of Palo Alto Networks firewalls that allows administrators to configure and enforce firewall policies based on users and user groups, instead of or in addition to network zones and addresses.

User-ID identifies the user on the network and the IP addresses of the computers the user is logged into to effectively enforce firewall policies. User-ID can also retrieve user and group information from a connected LDAP directory, allowing administrators to configure policies based on user groups, which are then translated into a list of users.

### How User Identification Works

The functionality provided by User-ID requires the collection of information from the network and directory servers. The following elements are involved in the information collection:

- **Identifying users on the network**
  User-ID provides a variety of mechanisms to reliably identify network users and their associated login session information (computers and network addresses). Some of the mechanisms require the installation of a User-ID Agent on network servers to provide the

most transparent user experience. With PAN-OS 5.0 and greater, you can use the built in PAN-OS User Mapping feature that performs the same function as the User-ID Agent, but does not require an agent installation.

- **Event log monitoring**
  Whenever a user authenticates to the Active Directory (AD) domain, a Microsoft Windows server, or Microsoft Exchange server, an event log is produced. Users can be identified on the network by monitoring those servers for the corresponding login events.
- **Server session monitoring**
  Another method is to continually monitor servers for network sessions established by users on the network. When a user successfully authenticated to a server, the session table of the server provides the user name and network source the user is connecting from.
- **Client Probing**
  In a Microsoft Windows environment, the client system can provide information about logged on users through Windows Management Instrumentation (WMI) for authorized users and services. Probing Microsoft Windows clients on demand provides information on users logged into a client computer.
- **XML API**
  Other identification methods are not directly supported by the User-ID features and options. For these cases, an XML over SSL interface is available, allowing customized solutions to register valid users and their corresponding client address on the network with User-ID.
- **Captive Portal**
  If the user cannot be identified based on login information, an established session or client probe, the firewall can re-direct any outbound HTTP requests and re-direct the user to a web form. The web form can transparently authenticate the user through a NTLM challenge, which is automatically evaluated and answered by the web browser or through an explicit login page.
- **Shared computers**
  Shared computers, such as Microsoft Terminal Servers, are problematic for most implementations, because a number of users share the same system and therefore the same network address. In this case, an Agent can be installed on the Terminal Server, which then associates not just the network address, but also allocated port ranges to the logged in users.

## Identifying Users and Groups

Policy management on the basis of individual users is unmanageable; therefore, users need to be associated and tied to user groups. Every enterprise environment stores user information in a directory service, such as Microsoft Active Directory or Novell eDirectory. All of those directory services are accessible through LDAP or LDAP over SSL (LDAPS).

The directory services provide resolution of user names and the associated user groups, which allows firewall administrators to configure security policies for user groups rather than individual users.

# How User-ID Components Interact

The User-ID Agent, PAN-OS, and the Terminal Services Agent interact with each other to provide complete user identification services.

## User-ID Agent

The User-ID Agents identifies the user on the network using one or all of the mechanisms described previously in this chapter.

- **Gathering user and login information**
  The User-ID Agent can be configured to monitor up to 100 Microsoft Windows Servers for user login events. When the Agent first connects to a server, it automatically retrieves a list of the last login events from the domain controller. During normal operations, it continues to receive new event information. The User-ID Agent provides the collected information to the firewall to enforce policy based on users and groups.
- **Providing users and network addresses to connected devices**
  To provide user and network address information, the firewall establishes a persistent connection to the User-ID Agent and retrieves a list of all identified users and network addresses on its first connection and every hour. During each hour, the firewall retrieves changes that the Agent detects.
- **On demand user identification**
  If the firewall identifies a new network address in the network traffic for which no user is listed, it can contact the User-ID Agent and request it to identify the user. This is done through a WMI or NetBIOS probe to the specific network address. When the client identifies the user, a new network address and username association is created and provided to the firewall.

## PAN-OS User Mapping

The PAN-OS User Mapping feature performs the same functions as the User-ID Agent, but does not require an agent installation on domain servers. This feature communicates directly with Microsoft servers and Novell eDirectory servers to gather user to IP and user to group mapping information. The enumeration of the individual users in a user group is also performed on the firewall. LDAP search and NTLM authentication is also supported. Refer to "User Mapping Tab" on page 286 and "PAN-OS User Mapping Configuration" on page 291.

You can also configure the firewall to share the user mapping information with other PAN-OS firewalls in your network. This is useful if you want to designate one or more devices to perform user mapping functions and then share that information with other firewalls. By using the include and exclude lists, you can also filter certain networks from the collection process. Refer to "PAN-OS User Mapping Configuration" on page 291.

## Terminal Services Agent

The Terminal Services Agent (TS Agent) solves the problem of multiple users using the same machine at the same time, for example on a Microsoft Terminal Server. After it is installed on the server, it allocates specific port ranges to each individual user. Every user connection is established using a port within the specific allocated port range.

When a port range is allocated for a particular user, the Terminal Services Agent notifies every connected firewall about the allocated port range so that policy can be enforced based on user and user groups.

### PAN-OS LDAP Group Query

In addition to enforcing policy based on individual users, the firewall can also be configured to allow or block traffic for groups of users. The enumeration of the individual users in a user group is performed by the firewall.

For this purpose, a LDAP server entry and group mapping settings need to be configured. The resulting LDAP query retrieves user groups and the corresponding list of group members.

This operation is performed every time a new configuration is submitted. Changes in group membership are detected through specific LDAP searches that retrieve only the groups and their member list that changed since the last search was performed.

## User Identification Agents

The User Identification Agent (User-ID Agent) is a Palo Alto Networks application that is installed on your network to obtain needed mapping information between IP addresses and network users. The User-ID Agent collects user-to-IP address mapping information from the domain controller security logs and provides it to the firewall for use in security policies and logs.

Refer to "PAN-OS User Mapping" on page 283, which performs the same functions as the User-ID Agent, but does not require an agent to be installed on the domain servers.

The IP address-to-user name mapping relies on the following mechanisms:

- For Active Directory, the security logs are continually monitored on the domain controller to detect user login events that contain user and IP address information.
- For Active Directory, a direct connection is required to all Domain Controllers to monitor user session activity and determine the user IP addresses.
- For eDirectory, when a user logs in, the IP address information is stored in eDirectory and retrieved by the User-ID Agent.
- The host PC is polled to verify IP address and user information using Windows Management Instrumentation (WMI) or Network Basic Input/Output System (NetBIOS). This occurs every 20 minutes to verify that the IP address-to-user name mapping is still correct and also when an IP address is seen that does not have an associated user name.
- The User-ID Agent application programming interface (API) is used to send information on user IP addresses to the User-ID Agent.
- User group mapping is performed through LDAP queries on directory servers. The firewall performs LDAP queries directly, but can use a configured User-ID agent as a LDAP proxy in cases where caching is desirable or direct access from the firewall to the directory server is not possible."

***Note:*** *User identification mapping requires that the firewall obtain the source IP address of the user before the IP address is translated with NAT. If multiple users appear to have the same source address, due to NAT or use of a proxy device, accurate user identification is not possible.*

In addition to the User-ID Agents, the firewall supports a Terminal Services Agent (TS Agent) that allows the firewall to identify individual users who are supported by the same terminal server. The firewall also supports captive portals for situations in which the User-ID Agent is unable to associate a user with an IP address.

Refer to the following sections for further information:

- "Captive Portals" in the next section
- "Configuring the Firewall for User Identification" on page 286
- "Setting Up the User-ID Agent" on page 296
- "Setting Up the Terminal Services Agent" on page 301

## Captive Portals

If the User-ID Agent is unable to associate a user with an IP address, a captive portal can take over and authenticate the user with a web form or NT LAN Manager (NTLM) challenge.

To receive the web form, users must be using a web browser and be in the process of connecting. Upon successful authentication, users are automatically directed to the originally requested web site. The firewall can now execute policies based on the user information for any applications passing through the firewall, not just for applications that use a web browser.

The following rules apply to captive portals:

- Captive portal rules work for HTTP and HTTPS web traffic.
- If the action for the rule is "web form," a web form is presented to the user to prompt for a password.
- If the rule is "NTLM" and the browser is Internet Explorer or Firefox, the firewall performs an NTLM authentication challenge (transparent to the user). If another browser is used, the web form is presented.

If the above-mentioned captive portal rules do not apply because the traffic is not HTTP or there is no rule match, then the firewall applies its IP-based security policies (as opposed to user-based security policies).

## Configuring the Firewall for User Identification

▶ *Device > User Identification*

Use the settings on this page to configure the firewall for user identification.

- **User Mapping tab**—Specify settings to use the PAN-OS User Mapping feature to provides accurate mappings between IP addresses and logged in users as well as user to group membership mapping. This option performs the same functions as the User-ID Agent directly from the firewall, so no agent is required on the domain controllers.
- **User-ID Agents tab**—Specify settings to support the user identification agent, which provides accurate mappings between IP addresses and logged in users.
- **Terminal Services tab**—Specify settings to support the terminal services agent. Refer to "Setting Up the Terminal Services Agent" on page 301.
- **Group Mappings Settings tab**—Specify settings to support mappings that associate users with user groups. User group mapping is performed by the firewall.
- **Captive Portal tab**—Specify settings to support use of a captive portal for user identification. Refer to "Captive Portals" on page 285.

**Table 129. User-ID Agent Settings**

| Field | Description |
|---|---|
| **User Mapping Tab** | |
| WMI Authentication tab | **User Name**—Specify the domain\username account that has permissions to perform WMI queries on client computers.<br>**Password/Confirm Password**—Specify the account password. |
| Server Monitor tab | **Enable Security Log**—Select the check box to enable security log monitoring on Windows servers. Security logs will be queried to locate user to IP mapping information on the servers specified in the **Server Monitoring** list.<br>**Server Log Monitor Frequency (sec)**—Specify how often the firewall will query Windows servers for user to IP mapping information (default 1 second, range 1-3600 seconds).<br>**Enable Session**—Select the check box to enable monitoring of user sessions on the servers specified in the **Server Monitoring** list. Each time a user connects to a server, a session is created and this information can also be used to identify the user IP address.<br>**Server Session Read Frequency (sec)**—Specify how often the firewall will query Windows server user sessions for user to IP mapping information (default 10 second, range 1-3600 seconds).<br>**Novell eDirectory Query Interval (sec)**—Specify how often the firewall will query Novell eDirectory servers for user to IP mapping information (default 30 second, range 1-3600 seconds). |

**Table 129. User-ID Agent Settings (Continued)**

| Field | Description |
|---|---|
| Client Probing tab | **Enable Probing**—Select this check box to enable WMI/NetBIOS probing to each client PC identified by the user mapping process. Probing will help ensure that the same user is still logged into the client PC in order to provide accurate user to IP information.<br>**Probe Interval (min)**—Specify the client PC probe interval (default 20 minutes, range 1-1440 minutes).<br>In large deployments, it is important to set the probe interval properly to allow time to probe each client that has been identified. Example, if you have 6,000 users and an interval of 10 minutes, it would require 10 WMI request a second from each client.<br>***Note:*** *For WMI polling to work effectively, the* ***User Mapping*** *profile must be configured with a domain administrator account, and each probed client PC must have a remote administration exception configured in the Windows firewall. For NetBIOS probing to work effectively, each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled.* |
| Cache tab | **Enable User Identification timeout**—Select this check box to enable a timeout value for user to IP mapping entries. When the timeout value is reached, the user to IP mapping will be cleared and a new mapping will be collected. This will ensure that the firewall has the most current information as users roam around and obtain new IP addresses.<br>**User Identification Timeout (min)**—Set the timeout value for user to IP mapping entries (default 45 minutes, range 1440). |
| NTLM tab | **Enable NTLM authentication processing—**Select this check box to enable NT LAN Manager (NTLM) authentication processing. When Captive Portal is set up with a web form to capture user to IP information, the client will be transparently authenticated through an NTLM challenge. With this option enabled, the firewall will collect this information from the NTLM domain.<br>If both the PAN-OS User-ID collector and the User-ID Agent is installed on domain controllers, NTLM responses will go directly to the domain controller. When the firewall is configured to share its User-ID information with other firewalls, it can serve NTLM requests coming from other PAN-OS firewalls, performing the function of the User-ID Agent.<br>**NTLM Domain**—Enter the NTLM domain name.<br>**Admin User Name**—Enter the administrator account that has access to the NTLM domain.<br>**Password/Confirm Password**—Enter the password for the administrator account that has access to NTLM domain.<br>**Note:** You cannot enable this check box in more than one virtual system. |

**Table 129. User-ID Agent Settings (Continued)**

| Field | Description |
|---|---|
| Redistribution tab | **Collector Name**—Specify the collector name if you want this firewall to act as a user mapping redistribution point for other firewalls on your network.<br>The collector name and pre-shared key are used when configuring the User-ID Agents on the firewalls that will pull the user mapping information.<br>To enable a firewall to act as a re-distribution point, you also need to enable the<br>**User-ID Service** in **Network > Network Profiles > Interface Mgmt**.<br><br>**Pre-Shared Key/Confirm Pre-Shared Key**—Enter the pre-shared key that is used by other firewalls to establishing a secure connection for user mapping transfers. |
| Server Monitoring | Click **Add** to configure the domain server(s) that the firewall will contact to gather user IP mapping and user to group mapping information.<br>**Name**—Enter the domain server network address.<br>**Network Address**—Enter the network address of the domain server.<br>**Type**—Select the type of domain server (Microsoft Active Directory, Microsoft Exchange, or Novell eDirectory).<br>**Enable**—Select the check box to enable the domain server profile.<br>**Discover**—Click this option to discover Microsoft Active Directory servers via DNS. You can then select check box next to the servers you want to use to obtain user mapping information.<br>The firewall will discover domain controllers based on the domain name entered in the **Device > Setup > Management** tab **> General Settings** page **Domain** field. |
| Include/Exclude Networks | **Name**—Enter a name to identify the profile that will include or exclude a network for User-ID discovery purposes. This option allows you to include or exclude a network range for IP address-to-user name mapping. Example, if you exclude 10.1.1.0/24, User-ID will not try to find user names for IP addresses in the excluded range. This in turn will also include or exclude ranges for mappings sent to other PAN-OS firewalls.<br>When defining an include or exclude network range, an implicit exclude-all will be performed. For example, if you include 10.1.1.0/24, all other networks will be excluded. If you exclude 10.1.1.0/24, all networks will be excluded, so when using exclude you must also have an include network, otherwise all networks are excluded.<br>**Enabled**—Select this option to enable the include/exclude profile.<br>**Discovery**—Select the option to either **Include** or **Exclude** the defined network range.<br>**Network Address**—Enter a network range that you would like to include or exclude for IP address-to-user name mapping discovery. For example, 10.1.1.0/24. |
| **User-ID Agents Tab** | |
| Name | Enter a name to identify the User-ID Agent (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |

**Table 129. User-ID Agent Settings (Continued)**

| Field | Description |
|---|---|
| Virtual System | Select the virtual system from the drop-down list (if supported on the firewall model). |
| IP Address | Enter the IP address of the Windows PC on which the User-ID Agent is installed. |
| Port | Enter the port number on which the User-ID Agent service is configured on the remote host.<br>When using a PAN-OS firewall as a redistribution point for user mapping information and you configure the User-ID Agent to connect to that device, you must use port 5007. |
| Collector Name | Enter the collector name that the User-ID Agent will pull from to obtain user mapping information. Refer to the "User Mapping Tab" on page 286. |
| Collector Pre-shared Key/Confirm Collector Pre-shared key | Enter the pre-shared key that will be used to allow SSL connectivity between the User-ID Agent and the firewall that is acting as a distribution point for user mapping. |
| Use as LDAP Proxy | Select the check box if the User-ID Agent is to be used as a LDAP proxy instead of the firewall connecting directly to the directory service. |
| Use for NTLM Authentication | Select the check box to use the configured User-ID Agent to verify NTLM client authentication from the captive portal with the Active Directory domain. |
| Enabled | Select the check box to enable the user identification agent. |
| Refresh Connected | Click to refresh the connection status of User-ID Agent profiles. This button is located to the right of the Add and Delete buttons. |
| Custom Agent Sequence | This option allows you to define the sequence order in which the User-ID agent profiles will connect to the defined server. For example, if you have four agents identified in the sequence list, it will attempt to connect to the first agent listed, if that connection fails, it will connect to the next agent listed, and so on. If this option is not configured, the connection sequence will follow the order of the agents listed in the main page. |
| **Terminal Services Agent Tab** | |
| Name | Enter a name to identify the TS Agent (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Virtual system | Select the virtual system from the drop-down list (if supported on the firewall model). |
| Host | Enter the IP address of the Windows PC on which the TS Agent will be installed. You can also specify alternative IP addresses (see the last entry in this table). |
| Port | Enter the port number on which the User-ID Agent service is configured on the remote host. |
| Alternative IP Addresses | Enter additional IP addresses, if the server has multiple IP addresses that can appear as the source IP address for the outgoing traffic. |

**Table 129. User-ID Agent Settings (Continued)**

| Field | Description |
|---|---|
| **Group Mapping Settings Tab** | |
| Name | Enter a name to identify the user-to-group mapping for user identification (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Virtual system | Select the virtual system from the drop-down list (if supported on the firewall model). |
| Server Profile subtab | Specify the following settings:<br>• Select an LDAP server profile from the drop-down list, and specify the interval (seconds) after which the configuration is updated with any new server profile information.<br>• Group objects<br>– **Search Filter**—Specify an LDAP query that can be used to control which groups are retrieved and tracked.<br>– **Object Clas**s—Specify the definition of a group. For example, the default is objectClass=group, which means that the system retrieves all objects in the directory that match the group filter and have object-Class=group.<br>– **Group Name**—Enter the attribute that specifies the name of the group. For example in Active Directory, this attribute is "CN" (Common Name).<br>– **Group Member**—Specify the attribute that contains the members of this group. For example in Active Directory, this attribute is "member."<br>• User Objects<br>– **Search Filter**—Specify an LDAP query that can be used to control which users are retrieved and tracked.<br>– **Object Clas**s—Specify the definition of the a user object. For example in Active Directory, the objectClass is "user."<br>– **User Name**—Specify the attribute for user name. For example in Active Directory, this attribute is "samAccountName." |
| Group Include List | Locate groups in the **Available Groups** list. Click the ⊞ icon to add the groups to the **Included** list and click the ⊟ icon to remove groups from the list. |
| **Captive Portal Settings Tab** | |
| Enable Captive Portal | Select to enable the captive portal option for authentication. |
| Location | Select the virtual system from the drop-down list (if supported on the firewall model). |
| Idle Timer | Enter the length of time after which the captive portal page times out (1-1440 minutes, default 15 minutes). |
| Expiration | Specify the timeout interval (range 1 - 1440 minutes, default 60 minutes). |
| Redirect Host | Specify the hostname used for the HTTP redirect used to initiate the NTLM challenge sent to the client. |

**Table 129. User-ID Agent Settings (Continued)**

| Field | Description |
|---|---|
| Server Certificate | Select the HTTP SSL certificate used for captive portal.<br>**Note:** If you select **None**, the firewall will use the local default certificate to provide an SSL connection. |
| Certificate | Choose the certificate profile to use for client authentication. |
| Authentication Profile | Choose the profile to determine the authentication source for captive portal logins. |
| NTLM Authentication | For NTLM authentication, specify the following:<br>• **Attempts**—Specify the number of attempts after which the NTLM authentication fails.<br>• **Timeout**—Specify the number of seconds after which the NTLM authentication times out.<br>• **Reversion Time**—Specify the time after which the firewall will again try to contact the first agent in the list of User-ID Agents after the agent becomes unavailable.<br>***Note:*** *These options only apply to the User-ID Agent installed on domain servers and does not apply to the PAN-OS User-ID feature configured in the User Mapping tab on the User Identification page.* |
| Mode | Choose whether the captive portal will use redirection or be transparent to the user.<br>Redirection is required for NTLM and session cookie retention. With the redirection option, the firewall can set a cookie for future login requests. Future redirection then becomes transparent to the user if the browser has not been closed.<br>For session cookies, specify the following settings:<br>• **Enable**—Select the check box to configure an interval after which the redirection times out.<br>• **Timeout**—If **Enable** is selected, specify the timeout interval (range 60 - 10080 minutes, default 1440 minutes).<br>• **Roaming**—Select the check box if to retain the cookie if the IP address changes while the browser is open (for example, if the client moves from a wired to wireless network). The cookie is lost when the browser closes, whether or not **Roaming** is selected.<br>***Note:*** *To use the captive portal in redirect mode, you must enable response pages on the interface management profile assigned to the Layer 3 interface to which you are redirecting the active portal. Refer to "Defining Interface Management Profiles" on page 175 and "Configuring Layer 3 Interfaces" on page 130.* |

# PAN-OS User Mapping Configuration

This section describes the basic steps needed to configure a firewall to retrieve IP address-to-user name mapping data directly from domain servers. This feature does not require the installation of a User-ID Agent on the domain servers. The firewall can also be configured to redistribute the user mapping information to other firewalls.

Refer to the following sections:

- "Configuring PAN-OS User Mapping" in the next section

- "Configure a Firewall to Share User Mapping Data" on page 295

# Configuring PAN-OS User Mapping

This section describes the configuration information needed to gather IP address-to-user name mapping and group mapping data for User Identification.

In this example, there is no need to install a User-ID Agent client on the domain controller. The following components need to be configured:

- **Domain Account**—A domain account is needed to query Microsoft Active Domain Controllers to obtain user information from security logs and to perform LDAP queries.
- **User-ID Agent**—As of PAN-OS 5.0, you configure the User-ID agent using the **User Mapping** tab, which is used to configure the native version of the agent.
  In large environments where performing user mapping enumeration may start to cause performance issues on the firewall, you will want to use the agent installed on domain controllers, which is configured in the **User-ID Agents** tab. For information on installing this agent, refer to "Setting Up the User-ID Agent" on page 296.
- **LDAP Server Profile**—This profile is used to define the domain controller server information that the firewall will use to perform queries for gathering user and group mapping information.
- **Group Mappings**—These settings allow the firewall to associate users with user groups, which enables you to create policies based on directory groups. The LDAP Server Profile is used in the Group Mappings settings to define the servers that will be queried.
- **Policies**—You need to configure the appropriate security profiles so the firewall can communicate with the defined directory servers.

The following outlines the basic steps needed to configure PAN-OS User Mapping:

1. Obtain a login account and password that has permissions to the servers and hosts that you will use to retrieve user mapping information. The permissions needed is based on the methods you will enable to gather the user mapping information, which may be from

Microsoft Active Directory servers, Microsoft Exchange servers, Microsoft file servers, Novell eDirectory, or using WMI to access client PCs.

**Note:** In Windows 2008 or later domains, you can add an account to the "Event Log Readers" group to access event logs to obtain IP address-to-user name mapping information from event logs. In Windows 2003 domains, the account should be assigned "Manage Auditing and Security logs" permissions through group policy.

You can also add the account to Server Operators to grant permissions to view current open server sessions to obtain IP address-to-user name mapping information for Windows 2003 and 2008 domains.

For the LDAP group queries to identify user to group mappings, no special membership is needed, most accounts can obtain group membership in Active Directory.

WMI is reliable and is secured by NTLM or Kerberos based authentication. To perform WMI queries, an account is needed that has rights to read CIMV2 namespace on the client system. By default, domain administrators have this permissions.

2. Navigate to **Device > User Identification** and click the **User Mapping** tab.
3. Click the **Edit** icon on the upper right side of the **Palo Alto Networks User ID Agent Setup** window. On this page, you define the settings that will be used when collecting IP address-to-user name mapping information on the servers defined in the **Server Monitoring** list. Populate each tab based on your requirements. Refer to "User Mapping Tab" on page 286 for descriptions of each of the user mapping settings. The **Redistribution** tab is only used if you plan on using this firewall to distribute user mapping information to other devices.
4. Click **OK** to save the settings.
5. In the **Server Monitoring** section, click **Add** to add the domain servers that you will access to gather user information. Enter the server Name, Network Address, and Type (Microsoft Active Directory, Microsoft Exchange, or Novel eDirectory).
   You can also click the **Discover** icon to discover domain controllers using DNS lookups. The firewall will discover domain controllers based on the domain name entered in the Device > Setup > Management > General Settings page.
6. Select the check box next to each domain server that you will monitor.
7. To specify networks that will be included or excluded in the user mapping data, click the Add icon in the Include/Exclude section and define the networks you want to filter. This allows you to send filtered user mapping data to the other firewalls in your network that use this device to pull user mapping information.

***Note:*** *You can also exclude a list of user accounts from being mapped by creating an ignore-user list. When using the Palo Alto Networks User Mapping feature, the list of user accounts to exclude is created in the CLI by using the command* `set user-id-collector ignore-user <value>`.

*There is no limit to the number of accounts that can be added to this list.*

8. Click Commit to make the changes active.

   You can check the status of the directory server connection by running the `show user server-monitor state` or `show user server-monitor statistics` commands.

Now that the User-ID Agent is configured, you need to configure an LDAP profile.

1. Name your configuration and leave the **Administrator use only** check box unchecked.
2. List the directory servers that you want the firewall to use in the server list. You need to provide at least one server; two or more are recommended for failover purposes. The standard LDAP port for this configuration is 389.
3. Leave the "Domain" field empty, unless you want to configure multiple independent directories.
4. Select a directory "Type". Based on the selected directory type, the firewall can populate default values for attributes and objectclasses used for user and group objects in the directory server.
5. Enter the base of the LDAP directory in the "Base" field. For example, if your Active Directory Domain is "paloaltonetworks.local", your base would be "dc=paloaltonetworks, dc=local", unless you want to leverage an Active Directory Global Catalog. In the "Bind DN" field, enter a user name for a user with sufficient permission to read the LDAP tree. In an Active Directory environment, a valid username for this entry could be the "User Principal Name", e.g. "administrator@paloaltonetworks.local" but also the users distinguished name, e.g. "cn=Administrator,cn=Users,dc=paloaltonetworks,dc=local".
6. Enter and confirm the authentication password for the user account that you entered above.

If you want to create policies based on directory groups, you need to configure Group Mapping. This configuration defines how groups and users are retrieved from the directory and which users groups should be included for use in policies.

1. Navigate to **Device > User Identification**, click the **Group Mapping Settings** tab and then click **Add**.
2. In the **Server Profile** drop-down, select the LDAP profile that you created in the previous section. All LDAP attributes and ObjectClasses will be pre-populated based on the directory server type you selected in the LDAP Server Profile page.
3. Adjust the **Update Interval** field if needed. The default is 3600 seconds (1 hour).
4. In order to optimize LDAP queries and policy configurations you can specify a list of user groups you would like to include in policies. To do this, click the **Group Include List** tab.
5. The left pane will show **Available Groups**, either navigate to the desired group, or search for groups that you want to include in your queries. You can use the asterisks (*) to do pattern matching in your searches.
6. Select the desired group(s) and then click the symbol to add the group(s) to the **Included Groups** pane.
7. Click **OK** to save group mapping.

# Configure a Firewall to Share User Mapping Data

Once you configure a firewall to retrieve user mapping information as described in "Configuring PAN-OS User Mapping" on page 292, you can designate the firewall to provide this information to other PAN-OS firewalls on your network.

The illustration in Figure 35 shows that a redistribution firewall communicates with the domain server to retrieve user to IP mapping information. The redistribution firewall can use any method to collect user mapping and group mapping information as described in "How User Identification Works" on page 281, and will then act as a User-ID Agent to share this information with other firewalls. Firewalls FW1, FW2, and so on, are configured to pull the mapping information directly from the redistribution firewall and do not need to communicate directly with any domain servers.

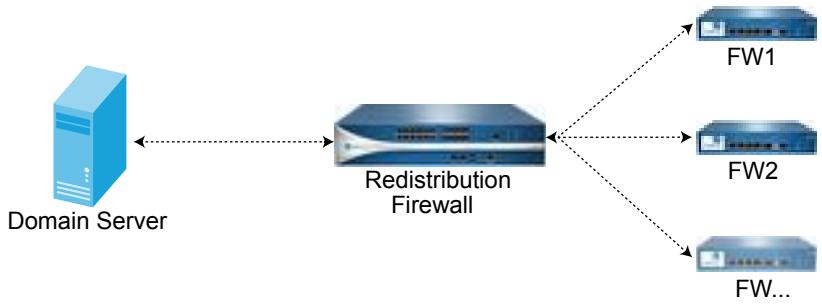


**Figure 35. Firewall configured to share user mapping information**

1. Navigate to **Device > User Identification** and click the **User Mapping** tab.
2. Click the **Edit** icon on the upper right side of the **Palo Alto Networks User ID Agent Setup** window and then click the **Redistribution** tab.
3. Enter a **Collector Name** and a **Pre-shared Key**. This information will be used by the firewall(s) that will pull user mapping information.
4. Click **OK** to save your changes.
5. Navigate to **Network > Network Profiles > Interface Mgmt** to enable the User-ID Agent Service.
6. Either **Add** a new Interface Management Profile, or modify an existing profile. Select the **User-ID Service** check box to enable the service.

The firewall is now ready to redistribute user mapping information to other firewalls.

To configure a firewall to retrieve information from a user mapping redistribution firewall, perform the following steps:

1. Navigate to **Device > User Identification** and click the **User-ID Agents** tab. Refer to "User-ID Agents Tab" on page 288 for setting details.
2. Click **Add** and enter a **Name** for the User Identification Agent profile and then populate the following fields:
   a. **Host**—Enter the host name or IP address of the firewall that is configured as a user mapping redistribution point.

b. **Port**—Enter 5007 for the port number. This is the designated port used for this service and cannot be modified.

c. **Collector**—Enter the collector name of the firewall that is configured as a user mapping redistribution point.

d. **Collector Pre-Shared Key**—Enter the collector pre-shared key of the firewall that is configured as a user mapping redistribution point.

e. **Use as LDAP Proxy**—Select this check box to use the firewall as an LDAP Proxy, instead of connecting directly to the directory server. This is used to query directory group and group membership information.

f. **Use for NTLM Authentication**—Select this check box to use the firewall to verify NTLM client authentication from the captive portal with the Active Directory domain.

g. Select the **Enabled** check box and click **OK** to save your settings and then **Commit** the configuration.

3. In the User-ID Agent window, look at the **Connected** column for the new User-ID Agent configuration. A green icon indicates that communication has been established with the firewall that is serving user mapping data.

   If a red icon appears, check traffic logs to determine why communications is not working properly. You can also check to see if any user mapping data has been received by running the operational command **show user ip-user-mapping** to view user mapping information on the dataplane, or **show user ip-user-mapping-mp** to view mappings on the management plane.

**Note:** When redistributing user mapping information to systems with multiple virtual systems, the following configuration is needed:

- Distribution Firewall—Each virtual system must have the collector configured with a pre-shared key. Refer to "Redistribution tab" on page 288. You also need to allow the User-ID Agent services from Network > Network Profiles > Interface Mgmt.
- Receiving Firewall—The collector name and pre-shared secret is needed in order to decide which virtual system User-ID information should be received.

# Setting Up the User-ID Agent

The User-ID Agent interfaces with Active Directory or eDirectory to communicate IP address-to-user name mapping information to the firewall.

The User-ID Agent is available for download from Palo Alto Networks. You can install the agent on one or more Windows PCs on your network to obtain user-specific information. When user identification is configured, the firewall's Application Command Center, App-Scope, and logs all include the user name in addition to the user IP address.

Follow the instructions in this section to install and configure the User-ID Agent.

***Note:*** *If the multiple virtual system capability is on, you can configure one or more agents per virtual system. This is useful to separate user identification in support of ISPs or other entities that maintain separate user records.*

## Installing the User-ID Agent

The system on which the User-ID Agent is installed must be running one of the following operating systems/directory server versions:

***Note:*** *The 32-bit or 64-bit versions of the below operating systems are supported.*

### Host Computers

- Microsoft Windows XP/Vista/7
- Microsoft Windows Server 2003/2008

### Directory Servers

- Microsoft Exchange
  - 2003 (6.5)
  - 2007 (8.0)
  - 2010 (14.0)
- Microsoft Active Directory
  - 2003
  - 2003r2
  - 2008
  - 2008r2
- Novell eDirectory Server 8.8

***Note:*** *Make sure that you choose the correct installation option for your client operating system (32-bit or 64-bit).*

Each PC that is included for user identification must be part of the authentication domain. For machines that are not part of the domain, you can use the captive portal capability to screen users and verify user names and passwords.

Refer to these sections for additional information:

- "Configuring the Firewall for User Identification" on page 286—Describes how to set up the firewall to communicate with the User-ID Agents and support captive portals.
- "Captive Portal Policies" on page 206—Describes how to set up captive portal policies.

To install the User-ID Agent, open the installer file and follow the on-screen instructions.

## Configuring the User-ID Agent

To open the User-ID Agent:

1. Choose **Start > All Programs > Palo Alto Networks > User-ID Agent**.

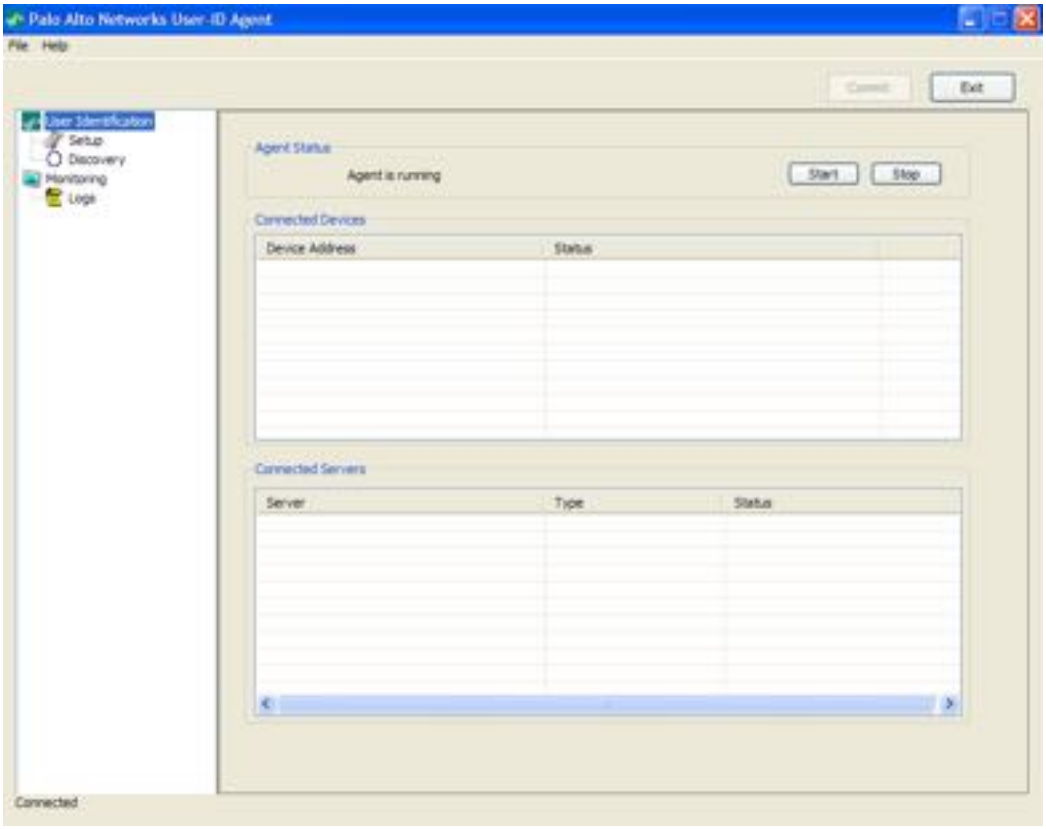


**Figure 36. User-ID Agent Window**

The window contains the following areas and functions:

- **Agent Status**—Displays the current status of the User-ID Agent.
- **Connected Devices**—Displays the list of devices that the User-ID Agent is currently connected to with associated status.
- **Connected Servers**—Displays the list of servers that the User-ID Agent is currently connected to with associated type and status.

To configure the User-ID Agent:

1. Choose **Start > All Programs > Palo Alto Networks > User Identification Agent**.
2. Click **Setup** to open the configuration window.

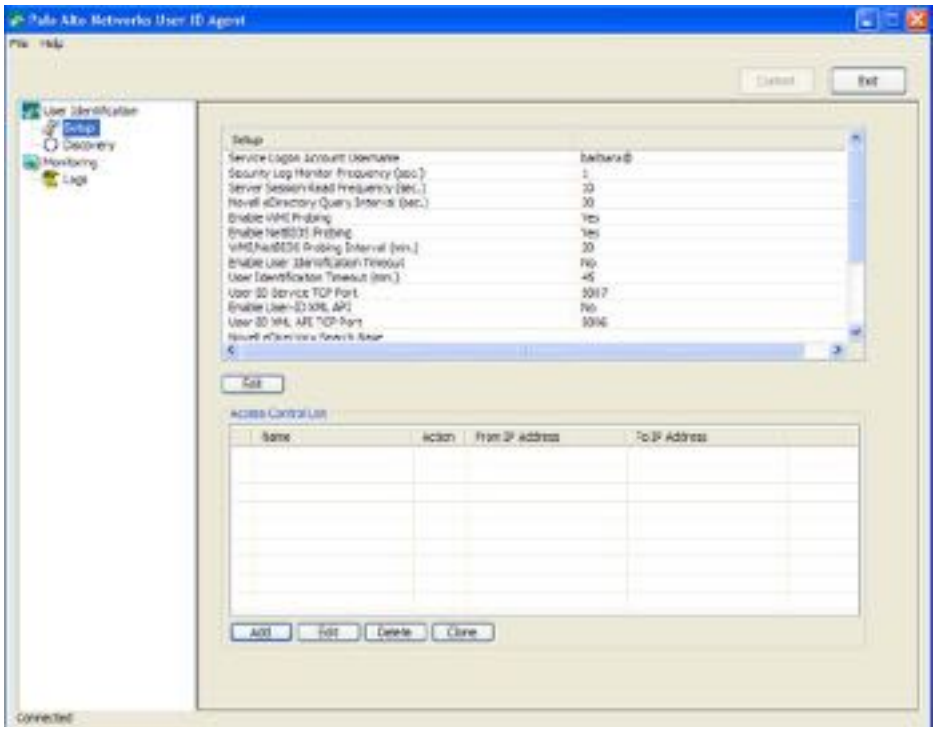

**Figure 37. User Identification Configuration Window**

3. The top of the window lists the current configuration settings. To modify the settings, click **Edit** just below the configuration summary and specify the following settings:
   - **Authentication**—Specify the user name and password to authenticate for Active Directory, WMI, NetBIOS, or eDirectory.
   - **Server Monitor**—Specify the frequency in seconds (default 1 second) for security log monitor and server session read (default 10 seconds) for Windows server and the query interval for Novell eDirectory query interval. (default 30 seconds).
   - **Client Probing**—Select the **Enable WMI Probing** check box if you want to enable WMI probing for each workstation and the **Enable NetBIOS Probing** check box if you want to enable NetBIOS probing for each workstation. Specify an interval between probes (seconds, default 20). An interval of 0 disables this feature.

***Note:*** *For WMI polling to work effectively, the Pan Agent service must be configured with a domain administrator account, and each probed client PC must have a remote administration exception configured in the Windows firewall.*

***Note:*** *For NetBIOS probing to work effectively, each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled.*

- **Cache**—Select the check box to enable timeout for the user ID and group cache, and specify the interval (minutes) after which the timeout occurs. Default 45 minutes.
- **Agent Service**—Specify the TCP ports for the user ID service (default 5007) and the user ID XML API (default 5006). Select the check box to enable use of the API.
- **eDirectory**—Specify the following settings:
  - **Search Base**—Specify the starting point or root context for agent queries. Example: dc=domain1, dc=example, dc=com.
  - **Bind Distinguished Name**—Specify the account to bind to the LDAP server. Example: cn=admin, ou=IT, dc=domain1, dc=example, dc=com.
  - **Bind Password**—Specify the bind account password. The agent saves the encrypted password in the configuration file.
  - **Search Filter**—Specify the search query for LDAP entries (default is objectClass=Person).
  - **Search Interval**—Specify the time interval between consecutive queries from the User-ID Agent (range 1-36000 secs, default 30 secs).
  - **Server Domain Prefix**—Specify a prefix to uniquely identify the user. Use if there are overlapping name spaces. Example: Different users with the same name from two different directories.
  - **Use SSL**—Select the check box to use SSL for eDirectory binding. If SSL is not selected, a pop-up window warns that clear text will be used for the login account and password.
  - **Verify Server Certificate**—Select the check box to verify the eDirectory server certificate when using SSL. Select the **Enable Group Cache** check box to enable the user-group membership cache. When this check box is selected, the user-group membership is cached, and when started, it first reloads the user-group membership from the cache to speed up the restart process.

4. Click **Save** to save the configuration.

   The User-ID Agent is restarted if the configuration is saved successfully. You can also click the **OK** button to save the configuration and restart the User-ID Agent. If you do not want to restart the User-ID Agent, click **Cancel** to close the dialog box.

***Note:*** *You can exclude a list of user accounts from being mapped by creating an ignore-user list. To do this, you create an ignore_user_list.txt file that is placed in the User-ID Agent folder on the domain server where the agent is installed.*

*As of PAN-OS 4.1 and later, there is no limit to the number of accounts that can be added to these lists. In PAN-OS 4.0 and earlier, the limit was 100 users.*

### Discovering Domain Controllers

The list of domain controllers available for domain login can be retrieved via DNS. To display the discovery options, click **Discover** on the side menu. You can perform the following tasks in this window:

- Add a configuration setting that allows an administrator to configure the User-ID agent to automatically discover available domain controllers for event log monitoring. Click **Add** or **Edit** in the **Servers** area and specify a server name, IP address, and type (Microsoft Active Directory, Microsoft Exchange, or Novell eDirectory).
- Specify an access control list for networks. Click **Add** or **Edit** in the **Include/exclude lists of configured networks** area, choose the include or exclude option, and specify the network name and address. You can also clone and then modify an existing entry.
- Click **Auto Discover** to automatically retrieve the list of available domain controllers from DNS and add those to the list of monitored servers.

### Monitoring User-ID Agent Operation

To view the list of currently discovered IP address-to-user name mapping information, click **Monitoring** on the side menu. You can search for users, or delete users from the list.

To view log entries for the User-ID agent, click **Logs** on the side menu. From this window you can search for log entries or clear the log.

### Uninstalling and Upgrading the User-ID Agent

To uninstall the User-ID Agent, open the Control Panel on the PC, select **Add or Remove Programs**, and remove the program **User Identification Agent**.

If you install a new version of the agent and the installer detects an existing installation on your PC, the installer automatically removes the older version before performing the installation.

We recommend that you back up the config.xml file before upgrading the User-ID Agent.

## Setting Up the Terminal Services Agent

The Terminal Server Agent (TS Agent) allows the firewall to support multiple users with the same source IP address by identifying the individual firewall users that the terminal server supports.

The TS Agent monitors the remote user sessions and reserves a different TCP/UDP source port range for each user. After a port range is allocated for the user, the TS Agent provides information to map the source port range to the user name.

In addition, the TS Agent requests that the TCP/UDP transport driver in the terminal server allocate the TS-Agent-specified source port instead of the operating system-determined ephemeral port for outbound TCP/UDP traffic. When the firewall receives the TCP/UDP traffic from the terminal server, it checks the source port and obtains the user ID in the ports-to-user map data for the terminal server.

For information on configuring the firewall for terminal services, refer to "Configuring the Firewall for User Identification" on page 286.

## Installing or Upgrading the Terminal Server Agent on the Terminal Server

You can install the TS Agent on the following platforms:

- Microsoft Terminal Services 2003
- Microsoft Terminal Services 2008
- Citrix Metaframe Presentation Server 4.0
- Citrix Metaframe Presentation Server 4.5, Citrix XenApp 5, 6

To install the TS Agent on the terminal server:

1. Download and open the installation file.
2. The installer first checks for platform compatibility. If the platform is not compatible, an error message is displayed.
3. The installer checks whether an existing TS Agent exists on the system. If the installer detects that the TS Agent already exists on the system (you are upgrading the TS Agent), it first uninstalls the agent before running the installer.
   - If you are installing a TS Agent that has a newer driver than the existing installation, the installation wizard prompts you to reboot the system after upgrading in order to use the new driver.
   - If you are installing a TS Agent with the same driver version as the existing installation, you can perform the installation as prompted, and do not need to reboot the system afterwards.
4. Follow the installer instructions to specify an installation location and complete the installation.

***Note:*** *If you specify a destination folder other than the default one, make sure that you use the same destination when you upgrade the TS Agent in the future. If you do not, the existing configuration will be lost and the default configuration will be used.*

5. Following installation, reboot the terminal server, if prompted to do so.

## Configuring the Terminal Server Agent on the Terminal Server

To configure the TS Agent on the terminal server:

1. Launch the TS Agent application from the **Start** menu.
2. The configuration panel opens with **Terminal Server Agent** highlighted on the left side of the window.

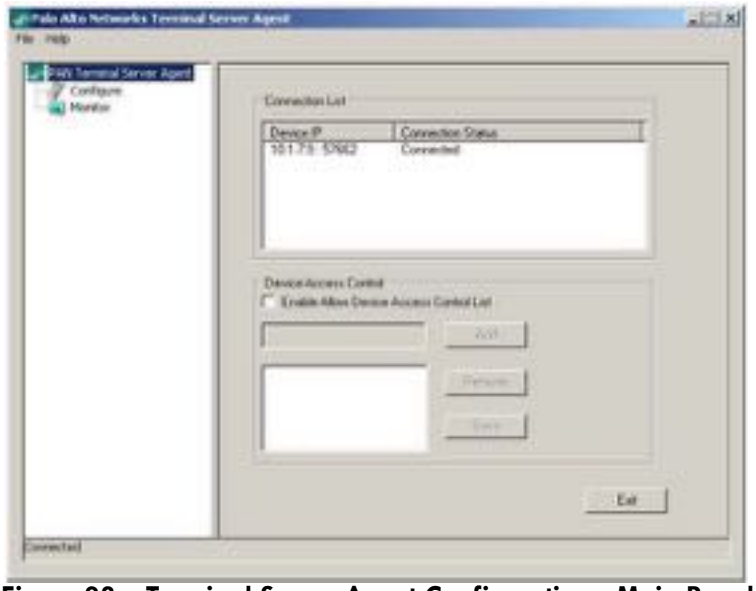


**Figure 38. Terminal Server Agent Configuration - Main Panel**

The connection list box shows all the Palo Alto Networks devices that connect to the TS Agent. The **Device IP** column shows the device IP and port; and the **Connection Status** column indicates whether the status is Connected, Disconnected, or Connecting. Disconnected items are removed from the **Connection List** box when you close and then reopen the TS Agent configuration window.

3. Select the **Enable Device Access Control List** check box if you want to explicitly list the firewalls that the TS Agent will accept. Add each device IP address and click **Add**. Click **Remove** to delete an address from the list. Click **Save** to save the allow list.

4. Click **Configure** to display the configuration settings.

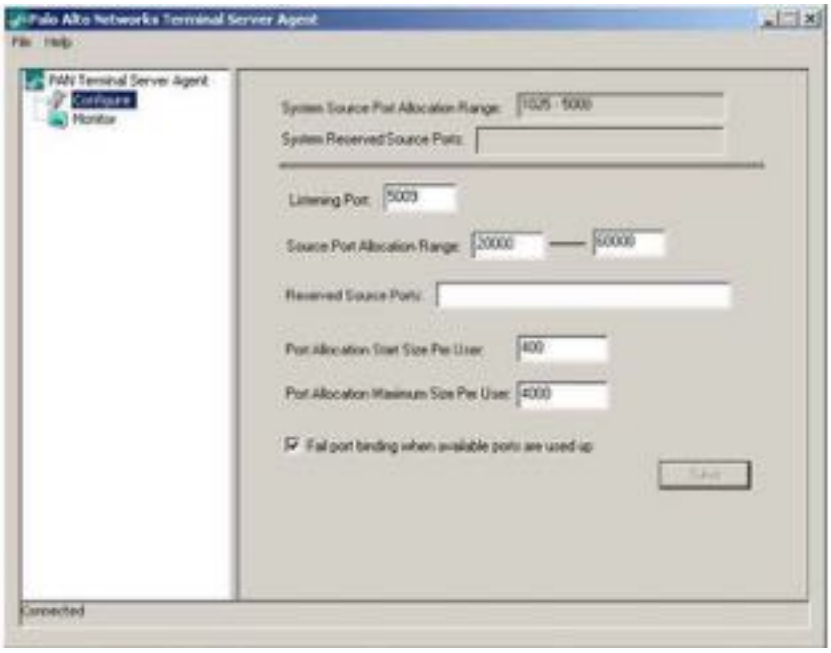


**Figure 39. Terminal Server Agent Configuration - Configure Panel**

5. Configure settings as described in the following table, and then click **Save**.

***Note:*** *If you enter an incorrect parameter and then attempt to save the configuration, a message is displayed to indicate that the configuration will not be saved unless you modify the parameter correctly.*

**Table 130. Terminal Server Agent Configuration Settings**

| Field | Description |
|---|---|
| System Source Port Allocation Range | Displays the port range for system processes that are not associated with individual users. When a server process opens a socket to send a UDP packet or set up a TCP connection, it must obtain a source port from the server operating system. The server automatically allocates a source port (an ephemeral port) for this process. Format is low-high (default 1025-5000).<br>The system port range must not overlap with the Source Port Allocation Range. If they overlap, an application using the system ephemeral source port range could mistakenly be identified as a particular user if the operation system allocated source port falls within the port range allocated for that user.<br>***Note:*** *Modifying this value requires a Registry change and cannot be done from this panel.* |
| System Reserved Source Ports | Displays the port or ports to be excluded from the operating system source port allocation (because other server processes may use them).<br>You can enter a range: *low-high* (no default).<br>***Note:*** *Modifying this value requires a Registry change and cannot be done from this panel.* |

**Table 130. Terminal Server Agent Configuration Settings (Continued)**

| Field | Description |
|---|---|
| Listening Port | Enter the port on which the terminal server will listen for communications from Palo Alto Networks firewalls (default 5009). |
| Source Port Allocation Range | Enter a port allocation range for user sessions.<br>This setting controls the source port allocation for processes belonging to remote users (default 20000-39999). If a port allocation request comes from system services that cannot be identified as a particular user process, the TS Agent lets the system allocate the source port from the system port range, excluding system reserved source ports.<br>***Note:*** *Make sure that this port range does not overlap with the System Source Port Allocation Range. If they overlap, an application using the system ephemeral source port range could mistakenly be identified as a particular user if the operation system allocated source port falls within the port range allocated for that user.* |
| Reserved Source Ports | Enter the reserved port allocation range for user sessions. These ports are unavailable for user sessions.<br>To include multiple ranges, use commas with no spaces, as in this example: 2000-3000,3500,4000-5000.<br>Format is *low-high* (no default). |
| Port Allocation Start Size Per User | Enter the number of ports that the TS Agent will first allocate when the remote user logs in (default 200).<br>When the remote user logs on, the TS Agent allocates a port range from the Source Port Allocation Range with this specified size. This allows identification of user traffic based on the source port. |
| Port Allocation Maximum Size Per User | Enter the maximum number of ports that the TS Agent can allocate for a remote user session (default 200).<br>If the **Port Allocation Start Size Per User** setting is not sufficient for the user session, the TS Agent will allocate additional ports up to this maximum. |
| Fail port binding when available ports are used up | Select the check box as appropriate:<br>• If the check box is selected (default), the port request from this user's application fails if the user application has used all available ports. As a result, the application may fail to send traffic.<br>• If the check box is not selected, the port request from this user's application is granted from the **System Source Port Allocation Range** even if the user application has used all the available ports. The application can send traffic; however, the user ID of the traffic is unknown. |

6. Click **Monitor** to display the port allocation information for all terminal server users.

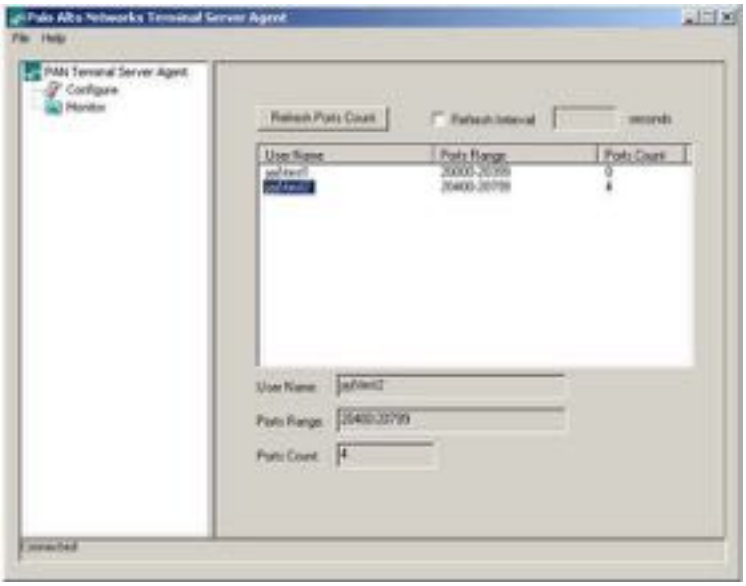


**Figure 40. Terminal Server Agent Configuration - Monitor Panel**

7. View the displayed information. For a description of the type of information displayed, refer to the following table.

**Table 131. Terminal Server Agent Monitor Information**

| Field | Description |
|---|---|
| User Name | Displays the user name. |
| Ports Range | Displays the current allocated source ports for this user. Multiple ranges are separated by commas (for example, "20400-20799, 20500-20599").<br>The size of the port ranges is limited by the "Port Allocation Start Size Per User" and "Port Allocation Maximum Size Per User" configuration parameters, as described in Table 130. |
| Ports Count | Indicates the number of ports in use. |

8. Click the **Refresh Ports Count** button to update the **Ports Count** field manually, or select the **Refresh Interval** check box and configure a refresh interval to update this field automatically.

The following table lists the menu options available in the TS Agent application window.

**Table 132. Terminal Server Agent Menu Options**

| Field | Description |
|---|---|
| Configure | Open the Configuration panel. |
| Monitor | Open the Monitor panel. |
| Restart Service | Restart the TS Agent service. This option is not normally required and is reserved for troubleshooting. |
| Show Logs | Display the troubleshooting log. |

**Table 132. Terminal Server Agent Menu Options (Continued)**

| Field | Description |
|---|---|
| Debug | Select debugging options (None, Error, Information, Debug, or Verbose). |
| Exit | Quit the TS Agent application. |
| Help | Display TS Agent version information. |

## Uninstalling the Terminal Server Agent on the Terminal Server

To uninstall the TS Agent, use the **Add/Remove Programs** control panel on the server. Remove the "Terminal Server Agent" application. You must reboot the system to complete the uninstallation.

## Chapter 8
# Configuring IPSec Tunnels

This chapter describes basic virtual private network (VPN) technology and provides details on configuring IP Security (IPSec) VPNs on Palo Alto Networks firewalls.

Refer to the following sections:


- "Virtual Private Networks" in the next section
- "IPSec and IKE" on page 311
- "Setting Up IPSec VPNs" on page 313
- "Sample VPN Configuration" on page 320
- "GlobalProtect Large Scale VPN Deployment" on page 323

# Virtual Private Networks

Virtual private networks (VPNs) allow systems to connect securely over a public network as if they were connecting over a local area network (LAN). The IP Security (IPSec) set of protocols is used to set up a secure tunnel for the VPN traffic, and the private information in the TCP/IP packets is encrypted when sent through the IPSec tunnel. The Palo Alto Networks firewall supports IPv4 tunnels as well as the option to enable IPv6 in the tunnel configuration. This option allows you to route IPv6 traffic over the IPv4 tunnel to provide confidentiality between IPv6 networks when IPv6 WAN connectivity is not available.

The Palo Alto Networks GlobalProtect Large Scale VPN feature provides a greatly simplified mechanism to roll out a hub and spoke VPN that utilizes certificates for authentication and automatically refreshes the authentication credentials periodically. Refer to "GlobalProtect Large Scale VPN Deployment" on page 323.

***Note:*** *In addition to IPSec VPNs, the firewall also supports Secure Sockets Layer (SSL) VPNs, which allow remote users to establish VPN connections through the firewall. Refer to* ***Chapter 9, "Configuring GlobalProtect"*** *for more information.*

The following figure shows a standard IPSec tunnel between two devices. The configuration can include a tunnel monitor on each side to alert the device administrator of tunnel failure and provide automatic failover. Tunnel monitors are useful if you want to be able to provide failover of IPSec traffic to another interface.

**Figure 41. IPSec Standard Configuration**

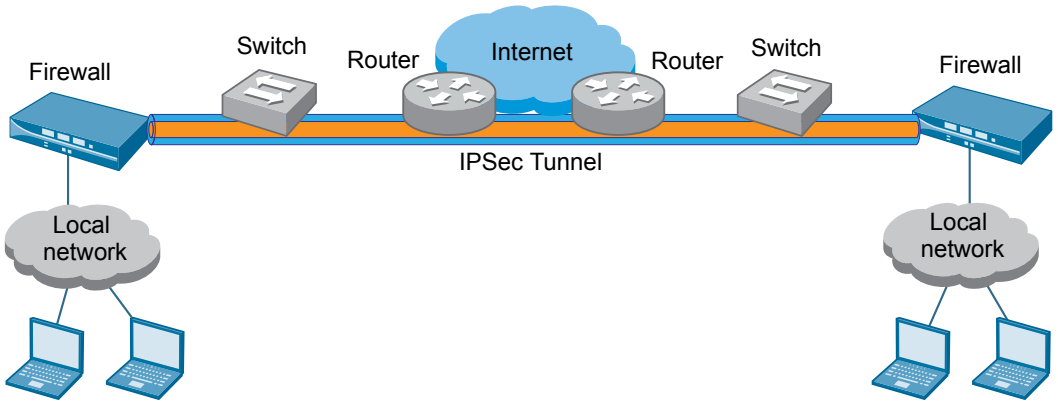

You can configure route-based VPNs to connect Palo Alto Networks firewalls at central and remote sites or to connect Palo Alto Networks firewalls with third-party security devices at other locations. With route-based VPNs, the firewall makes a routing decision based on the destination IP address. If traffic is routed to a specific destination through a VPN tunnel, then it is encrypted as VPN traffic. It is not necessary to define special rules or to make explicit reference to a VPN tunnel; routing and encryption decisions are determined only by the destination IP address.

The firewall can also interoperate with third-party policy-based VPN devices. To connect with a policy-based VPN, configure the Proxy ID for the tunnel. If multiple phase 2 tunnels are required, configure different Proxy IDs on each. Refer to "Setting Up IPSec Tunnels" on page 315.

For the IPSec connection between the firewalls, the full IP packet (header and payload) is embedded in another IP payload, and a new header is applied. The new header uses the IP address of the outgoing firewall interface as the source IP address and the incoming firewall interface at the far end of the tunnel as the destination IP address. When the packet reaches the firewall at the far end of the tunnel, the original packet is decrypted and sent to the actual destination host.

IPSec Security Associations (SAs) are defined at each end of the IPSec tunnel to apply all of the parameters that are required for secure transmission, including the security parameter index (SPI), security protocol, cryptographic keys, and the destination IP address. Encryption, data authentication, data integrity, and endpoint authentication are provided by IPSec SAs.

### VPN Tunnels

To set up VPNs, it is important to understand your network topology and be able to determine the required number of tunnels. For example:

- A single VPN tunnel may be sufficient for connection between a single central site and a remote site.
- Connections between a central site and multiple remote sites require VPN tunnels for each central - remote site pair.

Each tunnel is bound to a tunnel interface. It is necessary to assign the tunnel interface to the same virtual router as the incoming (clear text) traffic. In this way, when a packet comes to the firewall, the route lookup function can determine the appropriate tunnel to use. The tunnel interface appears to the system as a normal interface, and the existing routing infrastructure can be applied.

Each tunnel interface can have a maximum of 10 IPSec tunnels. This allows you to set up IPSec tunnels for individual networks that are all associated with the same tunnel interface on the firewall.

## IPSec and IKE

There are two ways to secure IPSec VPN tunnels:

- Configure the tunnel using manual security keys. This method is not recommended.
- Generate keys using Internet Key Exchange (IKE)

The same method must be applied to both ends of the IPSec tunnel. In the case of manual keys, the same key is entered at both ends; in the case of IKE, the same methods and attributes are applied at both ends.

IKE provides a standard mechanism for generating and maintaining security keys:

- **Identification**—The identification process involves recognition of the peers at both ends of the IPSec tunnel. Each peer is identified by IP address or peer ID (contained in the payload of the IP packet). The firewall or other security device at each end of the tunnel adds the identification of the peer at the other end into its local configuration.
- **Authentication**—There are two types of authentication methods: pre-shared key and PKI. Currently only the pre-shared key method is supported by Palo Alto Networks firewalls.

The firewall supports definition of IKE gateways, which specify the configuration information necessary to perform IKE protocol negotiation with peer gateways.

IKE configuration options include Diffie-Hellman Group for key agreement, Encryption algorithm, and hash for message authentication.

## IPSec and IKE Crypto Profiles

Crypto profiles are related to standard proposal fields in IKE negotiation.

- IKE Phase-1 authenticates the firewalls to each other and sets up a secure control channel. It uses the IKE-crypto profile for IKE SA negotiation.
- IKE Phase-2 is the negotiation, through the Phase 1 SA, of an actual tunnel for traffic between networks behind the respective firewalls. It uses the IPSec crypto profile for IPSec SA negotiation.

You can define IPSec and IKE crypto profiles that determine the protocols and algorithms used to negotiate the IPSec and IKE SAs.

Options for IKE SA:

- **Diffie-Hellman (DH) Group**—Select DH groups to use when generating public keys for IKE.
- **Encryption**—Select encryption algorithms.
- **Hash Algorithm**—Select hash algorithms.
- **Lifetime**—Specify the length of time that the negotiated key will stay effective.

Options for IPSec SA:

- **Encapsulating Security Payload (ESP)**—Select options for authentication, data integrity, confidentiality, and encryption.
- **Authentication Header (AH)**—Select options for authentication and data integrity. This option is not generally used.
- **Perfect Forward Security (PFS) Diffie-Hellman (DH) group**—Select DH groups to use in generating independent keys for IPSec.
- **Lifetime**—Specify the length of time that the negotiated key will stay effective.

For details on the specific protocols and algorithms supported for IPSec and IKE crypto profiles, refer to “Defining IKE Crypto Profiles” on page 318 and “Defining IPSec Crypto Profiles” on page 319.

# Setting Up IPSec VPNs

This section describes the multi-step process involved in setting up IPSec VPN tunnels. For detailed instructions, refer to the specified sections in this guide. For a sample configuration, refer to "Sample VPN Configuration" on page 320.

***Note:*** *Before you begin, make sure that your Ethernet interfaces, virtual routers, and zones are configured properly. Refer to "Firewall Interfaces" on page 127, "Virtual Routers and Routing Protocols" on page 153, and "Defining Security Zones" on page 151.*

To set up IPSec VPNs:

1. Plan the network topology and determine the required number of tunnels.
2. Define IKE gateways with the configuration information for IKE protocol negotiation with peer gateways. Refer to "Defining IKE Gateways" on page 314.
3. Configure the protocols and algorithms for identification, authentication, and encryption in VPN tunnels using IKE SA negotiation:
   - For IKEv1 Phase-1, refer to "Setting Up IPSec Tunnels" on page 315.
   - For IKEv1 Phase-2, refer to "Defining IPSec Crypto Profiles" on page 319.
4. Configure the parameters that are needed to establish IPSec VPN tunnels. Refer to "Setting Up IPSec Tunnels" on page 315.
5. Specify how the firewall will monitor the IPSec tunnels. Refer to "Viewing IPSec Tunnel Status on the Firewall" on page 319.
6. Set up static routes or assign routing protocols to redirect traffic into the newly established tunnels. The Border Gateway Protocol, Routing Information Protocol (RIP), and Open Shortest Path First (OSPF) options are supported; you can enable these protocols on the tunnel interface. Refer to "Virtual Routers and Routing Protocols" on page 153.
7. Set security policies to filter and inspect the traffic as described in "Security Policies" on page 187. Define the source and destination zones and specify the policy attributes as follows:
   - **Outgoing traffic entering the tunnel**—For source, use the clear text zone. For destination, use the tunnel interface zone.
   - **Incoming traffic egressing the tunnel**—For source, use the tunnel interface zone. For destination, use the clear text zone.

   After defining the zones for the security policy rule, set the source and destination addresses, applications, services, and security profiles to control access across the VPN tunnel.

***Note:*** *If the tunnel interface is in the same zone as the destination clear text traffic, the same policies will be used for VPN and clear text traffic. Ideally, you will want to put the tunnel interface in a separate zone, so tunneled traffic can use different policies.*

When these tasks are complete, the tunnel is ready for use. Traffic destined for the addresses defined for the tunnels is automatically routed properly and encrypted as VPN traffic based on the specific destination route added to the routing table.

***Note:*** *Without matching security rules, VPN traffic will be dropped by the firewall, when a security rule is required.*

*The IKE protocol will be triggered when necessary (for example, when traffic is routed to an IPSec tunnel with no keys or expired keys).*

*If there is a deny rule at the end of the security rulebase, intra-zone traffic is blocked unless otherwise allowed. Rules to allow IKE and IPSec applications must be explicitly included above the deny rule.*

## Defining IKE Gateways

▶ ***Network > Network Profiles > IKE Gateways***

Use the **IKE Gateways** page to define gateways that include the configuration information necessary to perform IKE protocol negotiation with peer gateways.

**Table 133. IKE Gateway Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the gateway (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Interface | Specify the outgoing firewall interface. |
| Local IP Address | Select the IP address for the local interface that is the endpoint of the tunnel. |
| Peer Type | Static IP address or dynamic option for the peer on the far end of the tunnel. |
| Peer IP Address | If the Static option is selected for peer type, specify the IP address for the peer on the far end of the tunnel. |
| Pre-Shared Key<br>Confirm Pre-Shared Key | Enter a security key to use for authentication across the tunnel. Applies for static and dynamic peer types. |
| ***Note:*** *The following advanced fields are visible if you select the check box for the* ***Show advanced Phase 1 options*** *link.* | |
| Local Identification | Choose from the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), KEYID (binary format ID string in HEX). If no value is specified, the local IP address will be used as the local identification value. |
| Peer Identification | Choose from the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), KEYID (binary format ID string in HEX). If no value is specified, the peer IP address will be used as the peer identification value. |
| Exchange Mode | Choose auto, aggressive, or main. |

**Table 133. IKE Gateway Settings (Continued)**

| Field | Description |
|---|---|
| IKE Crypto Profile | Select an existing profile or keep the default profile. |
| Enable Passive Mode | Select to have the firewall respond only to IKE connections and never initiate them. |
| Enable NAT Traversal | Select to have UDP encapsulation used on IKE and UDP protocols, enabling them to pass through intermediate NAT devices.<br>NAT traversal is used when NAT addressing is in place between the IPSec VPN terminating points. |
| Dead Peer Detection | Select the check box to enable and enter an interval (2 - 100 seconds) and delay before retrying (2 - 100 seconds). Dead peer detection identifies inactive or unavailable IKE peers through ICMP ping and can help restore resources that are lost when a peer is unavailable. |

***Note:*** *When a device is set to use the* ***auto*** *exchange mode, it can accept both main mode and aggressive mode negotiation requests; however, whenever possible, it initiates negotiation and allows exchanges in main mode.*

*You must configure the peer device with the matching exchange mode to allow it to accept negotiation requests initiated from the first device.*

## Setting Up IPSec Tunnels

▶ ***Network > IPSec Tunnels***

Use the **IPSec Tunnels** page to set up the parameters to establish IPSec VPN tunnels between firewalls.

**Table 134. IPSec Tunnel Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Name | Enter a name to identify the tunnel (up to 63 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.<br>The 63 character limit for this field includes the tunnel name in addition to the Proxy ID, which is separated by a colon character. |
| Tunnel Interface | Select an existing tunnel interface, or click **New Tunnel Interface** to create a new tunnel interface. For information on creating a tunnel interface, refer to "Configuring Tunnel Interfaces" on page 147. |
| Type | Select whether to use an automatically generated or manually entered security key. **Auto key** is recommended. |

**Table 134. IPSec Tunnel Settings (Continued)**

| Field | Description |
|---|---|
| Auto Key | If you choose **Auto Key**, specify the following:<br>• **IKE Gateway**—Refer to "Defining IKE Gateways" on page 314 for descriptions of the IKE gateway settings.<br>• **IPSec Crypto Profile**—Select an existing profile or keep the default profile. To define a new profile, click **New** and follow the instructions in "Defining IPSec Crypto Profiles" on page 319.<br>**Advanced**<br>• **Enable Replay Protection**—Select this option to protect against replay attacks.<br>• **Copy TOS Header**—Copy the (Type of Service) TOS header from the inner IP header to the outer IP header of the encapsulated packets in order to preserve the original TOS information.<br>• **Tunnel Monitor**—Select this option to alert the device administrator of tunnel failures and to provide automatic failover to another interface. Note that you need to assign an IP address to the tunnel interface for monitoring.<br>– Destination IP—Specify an IP address on the other side of the tunnel that the tunnel monitor will use to determine if the tunnel is working properly.<br>– Profile—Select an existing profile that will determine the actions that are taken if the tunnel fails. If the action specified in the monitor profile is wait-recover, the firewall will continue to use the tunnel interface in routing decisions as if the tunnel remained active. If the fail-over action is used, the firewall will disable the tunnel interface, thereby disabling any routes in the routing table that use the interface. For more information, see "Defining Monitor Profiles" on page 177. |
| Manual Key | If you choose **Manual Key**, specify the following:<br>• **Local SPI**—Specify the local security parameter index (SPI) for packet traversal from the local firewall to the peer. SPI is a hexadecimal index that is added to the header for IPSec tunneling to assist in differentiating between IPSec traffic flows.<br>• **Interface**—Select the interface that is the tunnel endpoint.<br>• **Local Address**—Select the IP address for the local interface that is the endpoint of the tunnel.<br>• **Remote SPI**—Specify the remote security parameter index (SPI) for packet traversal from the remote firewall to the peer.<br>• **Protocol**—Choose the protocol for traffic through the tunnel (**ESP** or **AH**).<br>• **Authentication**—Choose the authentication type for tunnel access (**SHA1**, **SHA256**, **SHA384**, **SHA512**, **MD5**, or **None**).<br>• **Key/Confirm Key**—Enter and confirm an authentication key.<br>• **Encryption**—Choose an encryption option for tunnel traffic (**3des**, **aes128**, **aes192**, **aes256**, aes128ccm16, or **null** [no encryption]).<br>• **Key/Confirm Key**—Enter and confirm an encryption key. |

**Table 134. IPSec Tunnel Settings (Continued)**

| Field | Description |
|---|---|
| GlobalProtect Satellite | If you choose **GlobalProtect Satellite**, specify the following:<br>For an overview of the concept of GlobalProtect Satellites, refer to "GlobalProtect Large Scale VPN Deployment" on page 323.<br>• **Name**—Enter a name to identify the tunnel (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.<br>• **Tunnel Interface**—Select an existing tunnel interface, or click New Tunnel Interface. For information on creating a tunnel interface, refer to "Configuring Tunnel Interfaces" on page 147.<br>• **Portal Address**—Enter the IP address of the GlobalProtect Portal.<br>• **Interface**—Select the interface from the drop-down that is the egress interface to reach the GlobalProtect Portal.<br>• **Local IP Address**—Enter the IP address of the egress interface that connects to the GlobalProtect Portal.<br>**Advanced Options**<br>• **Publish all static and connected routes to Gateway**—Select this option to publish all routes from the satellite device to the GlobalProtect Gateway in which this satellite is connected.<br>• **Subnet**—Click **Add** to manually add local subnets for the satellite location. If other satellites are using the same subnet information, you must NAT all traffic to the tunnel interface IP. Also, the satellite must not share routes in this case, so all routing will be done through the tunnel IP.<br>• **External Certificate Authority**—Select this option if you will use an external CA to manage certificates. Once you have your certificates generated, you will need to import them into the device and select the **Local Certificate** and the **Certificate Profile** to be used. |
| **Proxy ID Tab** | |
| Proxy ID | Click **Add** and enter a name to identify the proxy. |
| Local | Enter an IP address or subnet in the format *ip_address/mask* (for example, 10.1.2.1/24). |
| Remote | If required by the peer, enter an IP address or subnet in the format *ip_address/mask* (for example, 10.1.1.1/24). |
| Protocol | Specify the protocol and port numbers for the local and remote ports:<br>• **Number**—Specify the protocol number (used for interoperability with third-party devices).<br>• **Any**—Allow TCP and/or UDP traffic.<br>• **TCP**—Specify the local and remote TCP port numbers.<br>• **UDP**—Specify the local and remote UDP port numbers.<br>***Note:*** *Each configured proxy ID will count towards the IPSec VPN tunnel capacity of the firewall.* |

### Important items to consider when configuring IPSec VPNs

Keep the following in mind when configuring IPSec VPNs:

- There must be a route to the remote network that is being tunneled.

- Pre-shared keys may be entered incorrectly on one of the devices. Pre-shared keys must always match.
- Phase 1 negotiation mode (aggressive/main) may not match on the devices. The negotiation mode must always match.
- A common misconfiguration is to enable perfect forward secrecy on only one side. It must be enabled on both sides.
- If the dynamic routing protocols advertise routes to public IP addresses through the IPSec tunnel, the device establishing the tunnel may attempt phase 1 negotiation with the destination set to the public IP rather than the endpoint of the IPSec tunnel. As a result, the connection is never created and routing fails. To address this problem, ensure that only private IP addresses route through the tunnel and that no public IP addresses or default routes exist in the routing table that points to the tunnel.
- A Proxy ID may be improperly entered for the device at the far end of the IPSec tunnel. This can occur because some vendors generate a default Proxy ID for IPSec communications that is not easily identified by the end user.

## Defining IKE Crypto Profiles

▶ *Network > Network Profiles > IKE Crypto*

Use the **IKE Crypto Profiles** page to specify protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPSec SA negotiation (IKEv1 Phase-1). Refer to "Virtual Private Networks" on page 310 for more information.

To change the ordering in which an algorithm or group is listed, select the item and then click the **Move Up** or **Move Down** icon. The ordering determines the first choice when settings are negotiated with a remote peer. The setting at the top of the list is attempted first, continuing down the list until an attempt is successful.

**Table 135. IKE Crypto Profile Settings**

| Field | Description |
|---|---|
| DH Group | Specify the priority for Diffie-Hellman (DH) groups. Click **Add** and select groups. For highest security, select an item and then click the **Move Up** or **Move Down** icon to move the groups with higher numeric identifiers to the top of the list. For example, move **group14** above **group2**. |
| Authentication | Specify the priority for hash algorithms. Click **Add** and select algorithms (md5, sha1, sha256, sha384, or sha512). For highest security, use the arrows to move **sha1** to the top of the list. |
| Encryption | Select the check boxes for the desired Encapsulating Security Payload (ESP) authentication options. Click **Add** and select algorithms (aes256, aes192, aes128, or 3des). For highest security, select an item and then click the **Move Up** or **Move Down** icon to change the order to the following: **aes256**, **aes192**, **aes128**, **3des**. |
| Lifetime | Select units and enter the length of time that the negotiated key will stay effective. |

## Defining IPSec Crypto Profiles

▶ *Network > Network Profiles > IPSec Crypto*

Use the **IPSec Crypto Profiles** page to specify protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPSec SA negotiation (IKEv1 Phase-2). Refer to "Virtual Private Networks" on page 310 for more information.

**Table 136. IPSec Crypto Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| IPSec Protocol | Choose an option from the drop-down list.<br>**ESP:**<br>• Click **Add** under Encryption and select the desired ESP encryption algorithms. For highest security, use the arrows to change the order to the following: 3des, aes128, aes192, aes256, or aes128ccm16.<br>• Click **Add** under Authentication and select the desired ESP authentication algorithms (md5, sha1, sha256, sha384, sha512, or none).<br>**AH:**<br>• Click **Add** under Authentication and select the desired AH authentication algorithms (md5, sha1, sha256, sha384, or sha512). |
| DH Group | Select the DH group. For highest security, choose the group with the highest identifier. |
| Lifetime | Select units and enter the length of time that the negotiated key will stay effective. The default is 1 hour. |
| Lifesize | Select optional units and enter the amount of data that the key can use for encryption. |

To change the ordering in which an algorithm or group is listed, select an item and then click the **Move Up** or **Move Down** icon to change the order. The listed order determines the order in which the algorithms are applied and can affect tunnel performance.

## Viewing IPSec Tunnel Status on the Firewall

▶ *Network > IPSec Tunnels*

To view the status of currently defined IPSec VPN tunnels, open the **IPSec Tunnels** page. The following status information is reported on the page:

- **Tunnel Status (first status column)**—Green indicates an IPSec SA tunnel. Red indicates that IPSec SA is not available or has expired.
- **IKE Gateway Status**—Green indicates a valid IKE phase-1 SA. Red indicates that IKE phase-1 SA is not available or has expired.
- **Tunnel Interface Status**—Green indicates that the tunnel interface is up (because tunnel monitor is disabled, or because tunnel monitor status is UP). Red indicates that the tunnel interface is down, because the tunnel monitor is enabled and the status is down.

# Sample VPN Configuration

This section describes a sample VPN configuration. In this sample, a branch office is connected with a headquarters office and branch office users are allowed to access a central server farm.

Refer to the following topics:

- "Existing Topology" in the next section
- "New Topology" on page 321
- "Configure the VPN Connection" on page 321
- "VPN Connectivity Troubleshooting" on page 322

## Existing Topology

Headquarters:

- Firewall public IP 61.1.1.1, on interface ethernet1/1, which is in zone "ISP", virtual-router "HQ"
- Server farm network is 10.100.0.0/16, connected through interface ethernet1/5 (IP 10.100.0.1), which is on zone "server", virtual-router "HQ"

Branch office:

- Firewall public IP is 202.101.1.1, on interface ethernet1/2, which is in zone "ISP-branch", virtual-router "branch"
- A PC network of 192.168.20.0/24, connected through interface ethernet1/10, which is on zone "branch-office", virtual-router "branch" (same as ethernet1/2)
- Security policy to allow traffic from zone "branch-office" to zone "ISP-branch" for Internet access from the PC network

The following figure shows the existing topology.

**Figure 42. Sample VPN Configuration - Existing Topology**

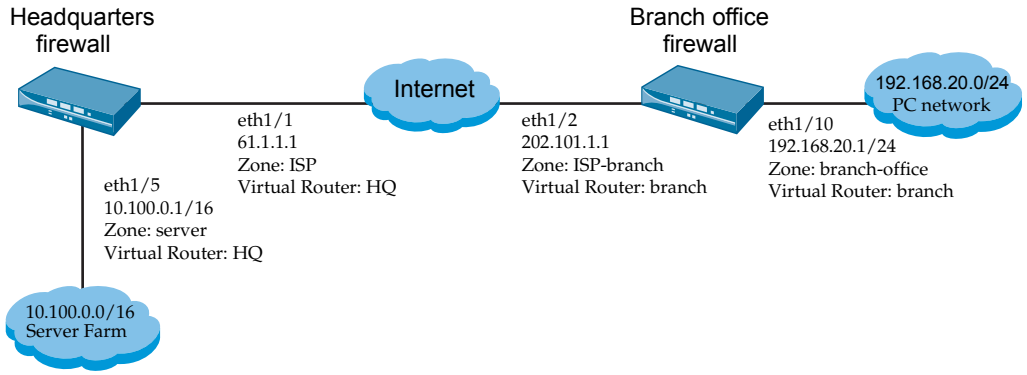

## New Topology

Headquarters:

- Create a new security zone “branch-vpn.”
- Add a tunnel interface tunnel.1 to zone “branch-vpn” and assign an IP address from a private range (for example, 172.254.254.1/24)
- Add a static route to direct traffic to 192.168.20.0/24 (the branch office network) to the tunnel interface tunnel.1.
- Add a security policy to allow traffic from zone “branch-vpn” to zone “server.”

Branch office:

- Create a new security zone “central-vpn.”
- Add a tunnel interface tunnel.2 to zone “central-vpn” and assign an IP address from private range (for example, 172.254.254.20/24).
- Add a static route to direct traffic to 10.100.0.0/16 (the server farm network) to the tunnel interface tunnel.2.
- Add a security policy to allow traffic from zone “branch-office” to zone “central-vpn.”

The following figure shows the tunnel information for the new topology.

**Figure 43. Sample VPN Configuration - new Tunnel Information**

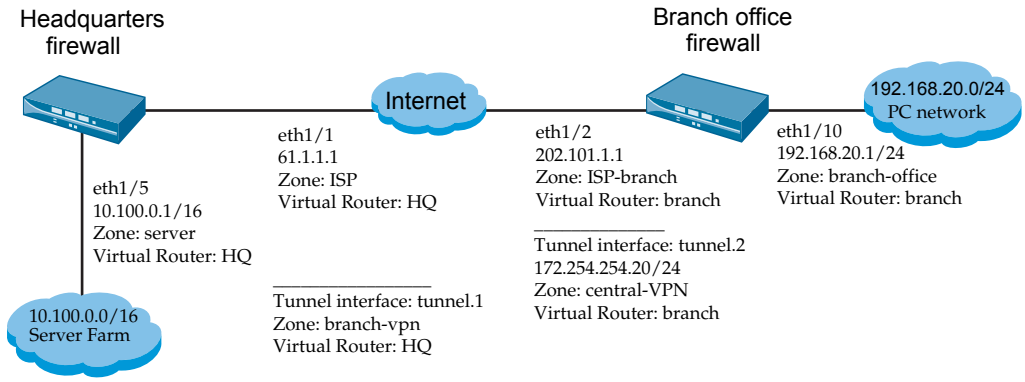


## Configure the VPN Connection

Headquarters:

- Create an IKE gateway “branch-1-gw” with these parameters:
  - Peer-address: **dynamic (or 202.101.1.1)**
  - Local-address: **ethernet1/1**
  - Peer-ID: type is FQDN: **branch1.my.domain**
  - Authentication: **pre-shared-key newvpn**
  - Protocol: keep default values

- Create an IPSec tunnel “branch-1-vpn” with these parameters:
  - ike-gateway-profile: **branch-1-gw**
  - ipsec-crypto-profile: leave as default
  - Tunnel interface: **bind with tunnel.1**
- On servers in the server farm, check the routing table and verify that the destination 192.168.20.0/24 is reachable through 10.100.0.1.

Branch office:

- Create an IKE gateway “central-gw” with these parameters:
  - Peer-address: **61.1.1.1**
  - Local-address: **ethernet1/2**
  - Local-ID: type is FQDN: **branch1.my.domain**
  - Authentication: pre-shared-key **newvpn**
  - Protocol: keep default values
- Create an IPSec tunnel “central-vpn” with these parameters:
  - ike-gateway-profile: **central-gw**
  - ipsec-crypto-profile: leave as default
  - Tunnel interface: bind with **tunnel.2**

Configuration Notes:

- If 202.101.1.1 is set as the peer-address parameter in “branch-1-gw” on the central site, setting the local-id and peer-id parameters becomes unnecessary (the field can be left empty). Note that treatment of these two parameters must be the same, because these two fields are matched during IKE negotiation.
- The proxy-id is left empty for route-based VPNs such as this.

After configuring the parameters and committing the configuration, the new VPN should work. If connectivity issues arise, refer to “VPN Connectivity Troubleshooting” in the next section.

## VPN Connectivity Troubleshooting

***Note:*** *The parameter values in this section refer to the sample configuration. Refer to “Configure the VPN Connection” on page 321.*

To troubleshoot issues regarding VPN connectivity:

1. Double check configurations on both sites.
2. Use the **ping** utility to verify connectivity between the central and branch offices (202.101.1.1 and 61.1.1.1). For this to work, there must be a management profile on the interface that allows **ping**.

3. Use the **ping** utility to verify connectivity between the server farm and the central firewall (ethernet1/5). For this to work, there must be a management profile on the interface that allows **ping**.
4. Use the **ping** utility to verify connectivity between the branch network and the branch firewall interface (ethernet1/10). For this to work, there must be a management profile on the interface that allows **ping**.
5. On the branch-office site, use the CLI commands **test vpn ike-sa gateway central-gw** and **show vpn ike-sa gateway central-gw** to verify that IKE phase-1 SA can be created from the branch office.
6. On the central site, use the CLI command **show vpn ike-sa gateway branch-1-gw** to verify that IKE phase-1 SA can be created from the branch office.
7. On the branch office site, use the CLI command **test vpn ipsec-sa tunnel central-vpn** and **show vpn ipsec-sa tunnel central-vpn** to verify that IKE phase-2 SA can be created from the branch office.
8. On the central site, use the CLI command **show vpn ipsec-sa tunnel branch-1-vpn** to verify that IKE phase-2 SA can be created from the branch office.
9. Check the server routing table in the server farm. The destination 192.168.20.0/24 must be reachable through the central firewall's ethernet1/5 interface IP address.
10. To check the route setting, run the **traceroute** command from any PC in the branch office network, where the destination is one of servers in the server farm.
11. Run the **ping** utility from any PC in the branch office network, where the destination is one of servers in the server farm. Check the encryption and decryption counters shown in the output of the **show vpn flow** CLI command. Verify that these counters are incrementing and that none of the error counters are incrementing.
12. Examine the detailed error messages for IKE negotiation in the syslog or use the **debug ike pcap** command to capture IKE packets in PCAP format.

# GlobalProtect Large Scale VPN Deployment

This section contains the following topics:


- "Overview" in the next section
- "Deploying a Large Scale VPN Network" on page 324
- "Dynamic Routing Protocols and Large Scale VPNs" on page 333
- "Backing up a GlobalProtect Portal" on page 334


## Overview

Deploying large scale VPN networks can be a very time consuming and complicated process. Challenges include planning, authentication requirements, tunnel configuration, route planning, and the provisioning and the decommissioning of components and devices. With

Palo Alto Networks firewalls, this process has been greatly simplified by taking advantage of the deployment and management methods used in the GlobalProtect VPN feature, which has previously only been used for remote access for client PCs.

The illustration in Figure 44 shows that each satellite connects to the portal to obtain a certificate for authentication and then downloads an initial VPN configuration. Once the initial configuration is completed, the satellite device establishes a VPN to all configured gateways defined on the portal using the same trusted certificate, creating a hub-and-spoke VPN network. Network and routing information is then shared between the gateway and satellite devices to create multiple paths to ensure that connectivity between the corporate and branch offices are always maintained.

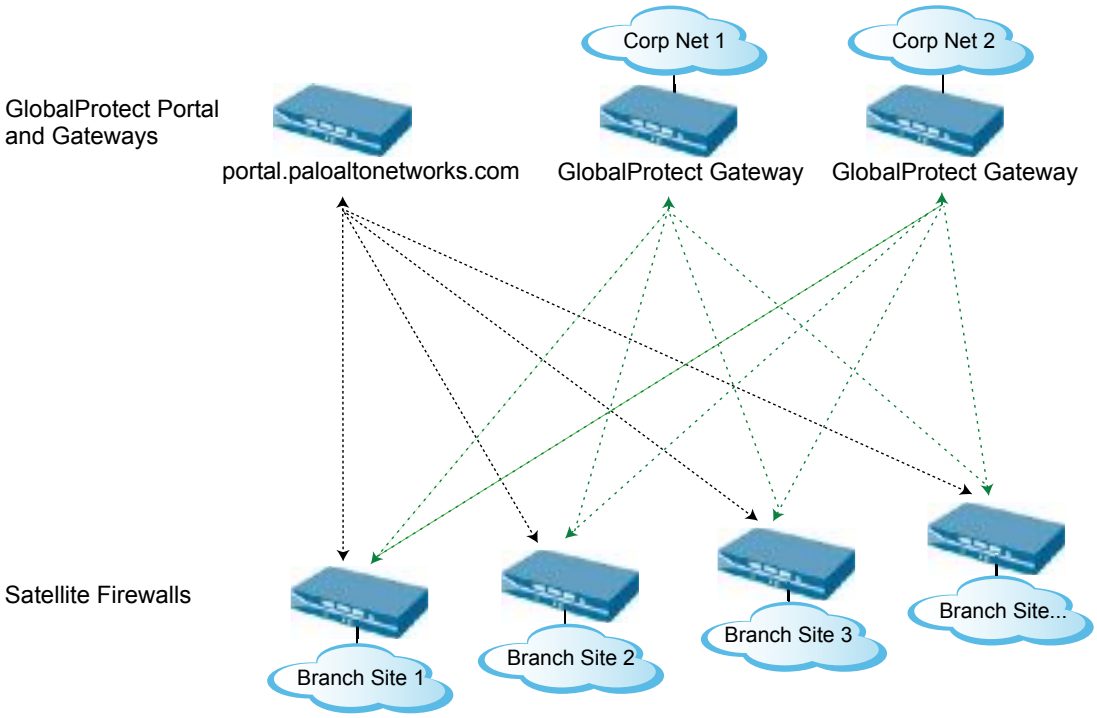


**Figure 44. Large Scale VPN Deployment**

# Deploying a Large Scale VPN Network

When deploying a large scale VPN network using Palo Alto Networks firewalls and GlobalProtect, you first configure the GlobalProtect Gateway(s), GlobalProtect Portal, and certificates that will be used for authentication of satellite devices. Once the portal and gateway devices are configured, you simply establish WAN connectivity from the satellite device, provision a tunnel interface, enter the host name of the Internet accessible GlobalProtect Portal interface and the satellite device configuration is initiated. The satellite device will then establish and maintain tunnels between all configured gateway devices and utilize tunnel monitoring to discover failures and initiate a failover. Once the portal and gateways are configured, you can easily add additional satellite devices.

It is important to keep a couple things in mind when deploying multiple branch offices:

- If you have branch offices that will use duplicate subnets, the satellite device must NAT all traffic to a unique IP or subnet. One of the simplest options is to source NAT all satellite-initiated traffic to the tunnel interface IP. When specifying the tunnel interface as a source NAT translated address, the firewall will use the IP address assigned to the tunnel interface by the highest priority gateway. If multiple gateways service different networks, it is recommended that you provision a unique IP address or subnet for each satellite that will be routable from all gateway devices. When duplicate subnets are in use, be sure that the satellites are not permitted to share connected routes with the gateways, otherwise routing issues will occur.
- Satellites must use tunnel monitoring (utilizing a configuration provided by the gateways) in order to recognize when a gateway fails. This also ensures that tunnels are maintained when there is no traffic flowing to the gateway.

The following section describes the components and basic steps needed to deploy a large scale VPN network using GlobalProtect.


- "Certificates and the OCSP Responder" on page 325
- "Global Protect Gateway Configuration" on page 328
- "GlobalProtect Portal Configuration" on page 330
- "GlobalProtect Satellite Configuration" on page 332
- "Dynamic Routing Protocols and Large Scale VPNs" on page 333


## Certificates and the OCSP Responder

Certificates are used in the large scale VPN environment to provide a mutual authentication scheme for the satellite devices and to automate deployment to reduce the steps needed to prepare the satellite devices.

The GlobalProtect Portal is configured as a root CA that will sign certificates for authentication for the satellite and gateway devices. The gateway uses a certificate profile with the portal's root CA in order to allow satellites to authenticate to the gateways. To add a new satellite device to the VPN network, you simply add the device serial number to the portal configuration or create a username and password that the satellite administrator will use to start the initial configuration. On the satellite, you enter the portal address in the IPSec Tunnels page and the satellite will connect to the portal and will send a certificate signing request (CSR). The portal then signs and returns the client certificate to the satellite as well as the root CA certificate to allow communication.

In the case of the gateway certificates, the portal will create a server certificate that is imported to the gateway, and the portal's root CA certificate is defined in the certificate profile on the gateway. This will allow all satellites to authenticate to the defined gateways.

To manage certificates in the environment, you create an Online Certificate Status Protocol (OCSP) profile on the portal that will be used to manage revocation status for local certificates issued to the satellite and gateway devices. On each satellite and gateway device, you will configure the OCSP option to point to the portal.

***Note:*** *You can also use an external root certificate authority (CA) for the GlobalProtect VPN deployment.*

For information on certificates, refer to "Importing, Exporting and Generating Security Certificates" on page 86. For information on the OCSP responder, refer to "OCSP Responder" on page 90.

Configure the OCSP Responder:

The OCSP responder needs to be configured first since it will be used when creating certificates for the satellite and gateway devices.

1. Navigate to **Device > Certificate Management > OCSP Responder**.
2. Click **Add** and enter a name for the OCSP responder.
3. Enter the **Host Name** of the portal.

Configure an OCSP Responder Management Profile:

This profile is used to allow the satellite and gateway devices to connect to the portal for OCSP requests over HTTP.

1. Navigate to **Network > Network Profiles > Interface Mgmt**.
2. Click **Add** and name the profile.
3. Select the **HTTP OCSP** check box. You may also want to enable ping for testing purposes.
4. Click **OK** to save your changes.
5. Assign the new management profile to the ingress interface of the portal by navigating to **Network > Interfaces**.
6. Select the interface that will be used as the ingress for the VPN devices and click the **Advanced** tab.
7. From the **Other Info** tab and the **Management Profile** drop-down, select the OCSP management profile that you created.
8. Click **Commit** to activate your changes.

Generate a root Certificate Authority (CA) certificate:

The portal root CA will sign and issue certificates for the satellite and gateway devices to provide mutual authentication for devices that are part of the large scale VPN network. The gateway will also have a certificate profile with the portal's root CA to allow authentication from the satellite devices.

1. Navigate to **Device > Certificate Management > Certificates**.
2. Click **Generate** and name the certificate.

3. Enter the **Common Name**, which is the IP or FQDN that will appear on the certificate. The common name is the FQDN or IP address of the interface where satellites will connect.
4. Check the **Certificate Authority** box.
5. In the **OCSP Responder** field, enter the OCSP responder name that you created earlier.

Generate a server certificate for the portal:

This certificate is used to authenticate satellite devices and to allow OCSP queries from the satellites and gateways.

1. On the GlobalProtect Portal, navigate to **Device > Certificate Management > Certificates**.
2. Click **Generate** and name the certificate and enter a common name.
   **Important:** The common name is the FQDN or IP address of the interface used for the satellite connections and will be unique to the portal.
3. In the **Signed By** drop-down, select the root CA of the portal. Set the desired cryptographic settings and other attributes and then click **Generate.**

Generate and import the gateway certificates:

Each gateway that will participate in the VPN network needs to have a server certificate that is signed by the GlobalProtect Portal root CA and the portal root CA certificate must also be imported into each gateway. This will allow the gateways to communicate with the portal and allows satellites devices to connect to the gateway to establish VPN connectivity.
In the below steps, you can alternatively export the root CA from the portal to the gateway, and then generate the gateway certificate directly on the gateway using the imported root CA to sign the certificate.

On the GlobalProtect Portal, navigate to **Device > Certificate Management > Certificates**.

1. Click **Generate** and name the certificate and enter a common name.
   **Important:** The common name is the FQDN or IP address of the interface used for the satellite to connect to the gateway and will be unique for each gateway.
2. In the **Signed By** drop-down, select the root CA of the portal.
3. In the **OCSP Responder**, it is important that you enter the URL of the portal using the following format: **http://*IP address or FQDN*/CA/ocsp**, for example http://paloaltonetworks.com/CA/ocsp. Certificates issued by the portal will have this path automatically generated in the certificate.
4. Set the desired cryptographic settings and other attributes and then click **Generate**.
5. Select the certificate you just created and then click **Export**, select the File Format **Encrypted Private Key and Certificate (PKCS12)**. Enter and confirm a passphrase that will be used when you import the certificate.
6. Follow the same steps to export the root CA certificate.
7. On the GlobalProtect Gateway, navigate to **Device > Certificate Management > Certificates**.
8. Click **Import**, enter a **Certificate Name** and then click **Browse** and select the gateway server certificate you exported earlier.

9. Select the file format **Encrypted Private Key and Certificate (PKCS12)**. Enter and confirm the passphrase you used when you exported the certificate and click **OK**.
10. Do the same for the portal root CA.

Configure a certificate profile on the gateway:

The certificate profile is used to authenticate the satellite device to the gateway to establish a the VPN.

1. Navigate to **Device > Certificate Management > Certificate Profile**.
2. Click **Add** and name the profile. The **Username Field** and **Domain** are optional when only Palo Alto Networks devices will connect to the gateway. You will need to configure these options if you have third-party VPN clients connecting to the gateway, such as with iOS and Android devices.
3. In the **CA Certificates** window, click **Add** and in the **CA Certificate** drop-down, select the root CA certificate that you imported from the portal.
4. In the **Default OCSP URL** field, enter the URL of the OCSP responder using the IP or FQDN of the portal. It is important that you enter the URL of the portal using the following format: **http://*IP address or FQDN*/CA/ocsp**, for example http://paloaltonetworks.com/CA/ocsp.

   If you have another device that can also handle OCSP requests, you can add a certificate for that profile in the **OCSP Verify CA Certificate** drop-down.
5. Click **OK**.
6. In the **Certificate Profile** window, click the **Use OCSP** check box.
7. Click **OK** to save your changes.
8. Click **Commit** to active the certificate configuration.

## Global Protect Gateway Configuration

The branch office satellite device will connect to the GlobalProtect Gateway to establish VPN connectivity to corporate resources. The satellite will use each gateway that is defined in the GlobalProtect Portal satellite configuration in order to create a hub and spoke VPN network. The gateway devices are also able to maintain multiple tunnels on a single tunnel interface, so you don't need a separate tunnel interface for each satellite device that is deployed. Authentication from the satellite to the gateway is handled by the certificates signed by the portal CA. Once configured and the VPN is established, routing information is also dynamically shared between the satellite and gateway devices.

Configuring the GlobalProtect Gateway device:

1. You should have already generated and imported the gateway server certificate and the portal root CA certificate. Refer to "Certificates and the OCSP Responder" on page 325.
2. Create a tunnel interface that is a logical interface used to terminate VPN tunnels. You can use the pre-defined tunnel interface, or create a new one. The interface also needs to be bound to the virtual router and placed in a security zone.
   It is recommended that you create a new zone, so you can control policies between the branch office and the zone that contains your protected resources. You could put the tunnel interface in the same zone as the protected resources, which will give the satellite

devices access to the resources, but that will not give you granular policy control for the satellite networks.
Refer to "Configuring Tunnel Interfaces" on page 147.

3. On the firewall that will be used as a GlobalProtect Gateway, navigate to **Network > GlobalProtect > Gateways** and click the **General** tab.
4. Click **Add** and name the gateway.
5. Select the interface with Internet connectivity that satellite devices will connect to in order to establish VPN connectivity and also select the **IP Address**.
6. Select the gateway server certificate that was signed by the portal and imported to the gateway.
7. Select the certificate profile in the **Certificate Profile** drop-down.
8. Click the **Satellite Configuration** subtab and check the **Tunnel Configuration** box.
9. Select the tunnel interface you created or select the default tunnel interface. You can also set **Replay attack detection** or **Copy TOS** (optional).
10. Set the **Configuration Refresh Interval**. This is the setting that each satellite will use to determine how often the satellite device should check the portal for configuration updates, such as the addition of a new gateway that the satellites can use. The default is 24 hours and the range is 1-48 hours.
11. Click the **Tunnel Monitoring** check box and enter a **Destination IP**. You can enter any desired IP address that is reachable, or if you leave this field blank, the monitoring profile will ping the gateway's tunnel interface IP address.
12. In the **Tunnel Monitor Profile** drop-down, configure a new profile and make sure the action is **failover**.
13. In the **IPSec Crypto Profile** you configure a new profile or select the default.
14. Click the **Network Settings** subtab and populate the DNS, IP Pool and Access route information that the satellite devices will use. The IP Pool is a range of IP addresses that will be assigned to the satellite device's tunnel interface that connects to the gateway when a VPN is established. The pool must be large enough to manage all satellites that will connect to the gateway.
    For the primary and secondary DNS fields, you can select **Inheritance Source** and select an existing interface that will be used to provide this information to the satellite devices.
15. Enter any desired **Access Route** information that you would like to have pushed down to the satellite devices. For example, you could provide access routes to implement split tunneling, so Internet traffic from the branch office does not pass through the VPN. If no access routes are provided, all traffic will pass through the tunnel.
16. Click the **Route Filter** tab and either leave this option unchecked to accept all routes sent by the satellite, or click **Accept published routes** and provide a list of subnets, which will cause all other routes to be filtered out. If you are using NAT, which will use the tunnel interface IP, you don't need to filter since no routes need to be published to the gateway.
17. Click **Commit** to activate the gateway configuration.

## GlobalProtect Portal Configuration

The portal is used to manage the VPN environment. It handles the enrollment process for new satellite devices, maintains the list of gateways to be used by satellites, and can acts as the root CA for the gateways and satellites. You can also use an external certificate authority (CA) if desired.

Once the satellite device receives its initial configuration and certificates from the portal, a VPN connection is established with all defined gateways, so it's not required to have a VPN between the portal and the gateways. The only time the gateway needs to communicate with the portal is for OCSP requests, which only requires HTTP access.
You can configure the portal as a gateway as well. This will allow you to have a gateway access point for satellite devices at the same location as the portal. You can then add additional external gateways for other locations throughout your global VPN network.


- "Create an Authentication Profile" on page 330
- "Security Zone and Interface" on page 331
- "Configuring the GlobalProtect Portal" on page 331


### Create an Authentication Profile

The authentication profile determines the authentication method that will be used by satellite devices to connect to the portal. When using this method, the satellite device administrator will enter a login and password when connecting to the portal for the first time. Once the device is authenticated, it will add the serial number and host name to the portal. You can use a Local DB account, RADIUS, LDAP using active directory, or Kerberos. In this case, the portal administrator does not need to add the satellite device serial number in the portal satellite configuration. Refer to "Authentication Profiles" on page 62.
You can also add the serial number of the satellite device as described in "Configuring the GlobalProtect Portal" on page 331. In both cases though, you must have an authentication profile defined in order to commit your configuration. If you manually add a device serial number to the portal configuration, the authentication profile is not used, but it is still required.

Configure the authentication profile:

1. To use a local account for authentication, on the portal navigate to **Device > Local Users Database > Users** and click **Add**.
2. Enter a username in the **Name** field and then enter a password in the Password/Confirm Password fields. You can also use a password hash by selecting the Password Hash button and pasting a password hash in the field.
3. Click **OK** to create the account.
4. Navigate to **Device > Authentication Profile** and click **Add**.
5. Enter a profile name in the **Name** field and in the **Authentication** drop-down, select **Local Database**. If you will be using RADIUS, LDAP, or Kerberos, select one of those options and configure the server. Once configured, you can then select usernames that are available on those systems.
6. In the **Allow List**, click **Add** and then select the local user you created.

### Security Zone and Interface

On each device that will participate in the VPN (Portal, Gateway, Satellite), you will need to configure an interface that will have access to the other devices in the VPN environment and put them in the desired security zone. This will allow each device to communicate over the WAN. It is recommended that you create a new zone, so you can control policies between the branch office and the zone that contains your protected resources.

The initial connection from the satellite to the portal for the certificate download and initial configuration will use SSL. Once the satellite is configured, it will establish an IPSec VPN with the configured gateways. Make sure that the security zones that you create have the **Enable User Identification** check box enabled, so users from the branch office can be identified if User-ID is configured. This is required for zones that will be used on the GlobalProtect Gateway. Refer to "Security Zones" on page 151.

### Configuring the GlobalProtect Portal

The portal configuration defines the network interface that will be used for satellite enrollment, authentication, and appearance of the portal. It can also act as a root CA and OCSP responder.

1. Navigate to **Network > GlobalProtect > Portals**.
2. Click **Add** and name the portal.
3. Select the interface that will act as the ingress interface for the satellite devices and then select the Internet accessible IP address from the drop-down.
4. Select the portal server certificate that you created earlier. Refer to "Certificates and the OCSP Responder" on page 325.
5. Select the **Authentication Profile** that you created earlier. The authentication profile is required, even though the satellite device will not use it if you add the satellite device serial number to the portal configuration.
6. In this case, a client certificate is not required on the portal since the portal automatically generates client certificates for satellite devices using the root CA.
7. Select a certificate profile to be used for satellite authentication. Alternatively, you can enter the satellite serial number in the **Satellite Configuration** sub tab, so the satellite administrator does not have to enter login credentials.
8. You can optionally configure appearance settings.
9. Click the **Satellite Configuration** tab that will define options for the satellite devices that will connect to the portal.
10. Click **Add** and name the **GlobalProtect Satellite** profile. This profile will define the satellite devices, enrollment method, and the list of gateways that the satellites will use in the VPN environment.
11. Set the **Configuration Refresh Interval**. This is the setting that each satellite will use to determine how often the satellite device should check the portal for configuration updates, such as new gateways being added.
12. In the **Devices** tab, you can manually add satellite device serial numbers, or the list will be updated when a satellite device connects to the portal for the first time using a login and password authentication.
    The next step provides information on the enrollment methods.

13. The portal can be configured to allow two types of authentication methods for the initial enrollment (you can configure both of the below options):

    – Manually add the satellite device serial number to the portal in the **Network > GlobalProtect > Portals > Satellite Configuration** page. Click **Add** > **Devices** then click **Add** again and enter the satellite device serial number. When the satellite device is configured at the branch office and establishes the initial connection to the portal, no authentication is required. Once the satellite device connects, the host name will be added to the portal configuration automatically.

    – Use a login name and password prompt from the satellite device by adding an enrollment user or group to the portal in the **Network > GlobalProtect > Portals > Satellite Configuration** page. Click **Add** > **Enrollment User/User Group** then click **Add** again and select a user or group from the drop down list or select **Any**.
    When using this method, you do not need to enter the satellite device information on the portal. When the device connects, the administrator configuring the satellite device will be prompted to login. Once authenticated, the satellites serial number and host name will automatically be added to the portal satellite list.

14. In the **Gateways** tab, enter the IP address or hostname of the gateway(s) that the defined list of satellite devices will use for VPN connectivity.
    You can also set the routing priority of the gateway from this page. If the satellite device has tunnels to multiple gateways, duplicate routes may exist. By setting a priority, the gateway with the highest priority will be used first. If the gateway fails, the next highest priority will be used. The priority range is from 1-25, since 25 is the limit on the number of gateways that a satellite can use.

15. Click **OK** to save your changes and then under **Trusted Root CA,** click **Add** and select the portal root CA.

16. In the **Issuing Certificate** drop-down, also select the portal root CA. This is the certificate that will be used to automatically generate certificates for the satellite devices during enrollment.

17. You can modify the validity period and certificate renewal period if desired. The validity period defines the lifetime of the certificate and the renewal period defines how often the certificate will be renewed. If you do not want your certificates to expire, just make sure the renewal period is lower than the validity period.

18. In the **OCSP Responder** drop-down, select the portal's OCSP responder that you created earlier and then click **OK** to save your changes.

19. Click **Commit** to activate your changes.

## GlobalProtect Satellite Configuration

On the branch office satellite device, you first configure an interface with WAN connectivity and set up a security zone and policy to allow the branch office LAN to communicate with the Internet. The satellite device is then configured with the host or IP address of the portal. Once the configuration is committed, the satellite will communicate with the GlobalProtect Portal, certificates are signed and downloaded with the initial VPN configuration and routing information. The satellite then connects to all configured GlobalProtect gateways to establish VPN connectivity from the branch to the corporate office. Refer to "GlobalProtect Satellite" on page 317 for field descriptions).

1. On the portal devices, you should have already entered the serial number of the satellite device you are configuring, or have created a username and login that can be used for the initial configuration. Refer to "Configuring the GlobalProtect Portal" on page 331.
2. Navigate to **Network** > **IPSec Tunnels** page.
3. Click **Add** to create a new IPSec Tunnel.
4. Enter a **Name** for the tunnel profile and select the default **Tunnel Interface** from the drop down, or create a new tunnel interface.
5. Choose **Type GlobalProtect Satellite** and enter the **Portal Address** (IP or hostname).
6. Select the **Interface** to be used to connect to the portal and then select the **Local IP Address** of the interface. In the example illustration shown in Figure 44 on page 324, the **Portal Address** would be "portal.paloaltonetworks.com".
7. Click the **Advanced** tab and configure the desired advanced options. It is recommended that you enable **Publish all static and connected routes to Gateway**, so that routes are shared between all devices in the VPN network.
   **Note:** If this branch office has the subnets as other branch offices in the VPN, do not enable this option unless you are using NAT.
8. Click **Commit** to activate the changes on the satellite device.

You can now test connectivity between the satellite devices and gateways. To view the status of the VPN on the satellite device, navigate to **Network > IPSec Tunnels**. You should see a green status indicator in the **Status** column.

# Dynamic Routing Protocols and Large Scale VPNs

When using dynamic routing protocols over your large scale VPN network, the following behavior should be noted in relation to flooding and replication:

- Routing Information Protocol (RIP)—Control packets are flooded to all tunnels (spokes) on the egress interface (hub to all spokes). Control packets are not replicated to remaining tunnels (spokes) when one tunnel receives a packet from the incoming spoke to the remaining spokes.
- OSPF—It is recommended that you use Point to Multipoint (P2MP). In broadcast mode, the replication behavior is the same as RIP.

Multicast replication behavior is the same as RIP and data packets are not replicated. Also, multicast distribution in a large scale VPN environment is supported where the multicast source resides behind the gateway and receivers reside behind the satellite VPN devices. Deployments in which the multicast source resides in a satellite location are not currently supported.

# Backing up a GlobalProtect Portal

In a large scale VPN configuration, the GlobalProtect Portal is a critical piece of the environment that manages the VPN configuration, all certificates used for authentication, and the list of satellite firewalls that are part of the VPN network. The certificate and satellite information is dynamic and changes often and is not part of the device configuration.

The "Export device state" feature is used to export the configuration and dynamic information from a firewall that is configured as a GlobalProtect Portal with the large scale VPN feature enabled. The export file contains a list of all satellite devices managed by the Portal, the running configuration at the time of the export, and all certificate information (Root CA, Server, and Satellite certificates). If the Portal experiences a failure, the export file can be imported to restore the Portal's dynamic information.

**Important:** You must manually run the device state export from **Device > Setup > Operations** by clicking the **Export device state** option, or you can create a scheduled XML API script to export the file to a remote server. This should be done on a regular basis since satellite certificates may change often.

To create the device state file from the CLI, run `save device state`.
The file will be named device_state_cfg.tgz and is stored in /opt/pancfg/mgmt/device-state.
The operational command to export the device state file is `scp export device-state` (you can also use tftp device-state export).

For information on using the XML API, refer to the document "PAN-OS XML-Based Rest API Usage Guide" at *https://live.paloaltonetworks.com/community/documentation.*

# Chapter 9
# Configuring GlobalProtect

This chapter describes GlobalProtect, which allows secure login from client systems located anywhere in the world:

- "Overview" in the next section
- "Setting Up GlobalProtect" on page 337
- "Setting Up and Activating the GlobalProtect Agent" on page 352

## Overview

GlobalProtect provides security for client systems, such as laptops, that are used in the field by allowing easy and secure login from anywhere in the world. With GlobalProtect, users are protected against threats even when they are not on the enterprise network by sending their traffic through a Palo Alto Networks firewall that is within close geographic proximity. The user's access level is determined by a host information profile (HIP) that notifies the firewall about the user's local configuration. The HIP information can be used for granular access control based on the security programs that are running on the host, registry values, and many other checks such as whether the host has disk encryption enabled.

The GlobalProtect agent can also be a Palo Alto Networks satellite device (firewall), but instead of agent software being downloaded, only the necessary certificate for authentication and VPN configuration is needed. For information on satellite devices, refer to "GlobalProtect Large Scale VPN Deployment" on page 323.

The following elements are used to provide GlobalProtect functionality:

- **Portal**—Palo Alto Networks firewall that provides centralized management for the GlobalProtect system.
- **Gateways**—Palo Alto Networks firewalls that provide security enforcement for traffic from GlobalProtect agents.
- **Agent**—Application that is installed on the client system and configured to connect to the portal and gateways to provide network access for user systems. The agent also provides information about the user's local configuration to the gateways.
- **Satellite**—A satellite device is a Palo Alto Networks firewall that is typically installed at a branch office and is initialized and managed by a GlobalProtect portal. After the satellite receives its initial configuration from the portal, it connects to all configured gateways to

establish VPN connectivity. By using GlobalProtect, firewalls can be quickly installed at branch offices with very little initial configuration. Refer to "GlobalProtect Large Scale VPN Deployment" on page 323.

The connection process works as follows:

1. The user downloads the GlobalProtect agent from the portal onto a client system, such as a laptop. The agent connects to the portal through an SSL connection and downloads a configuration file. This configuration file includes information about the various GlobalProtect gateways that can be used for connectivity.
2. The agent performs a reverse Domain Name System (DNS) lookup to determine whether the client system is currently on the internal enterprise network or on an external network.
3. If the connection is to the external network, the agent attempts to make SSL connections to all external gateways and then selects the best gateway.
4. Based on the configuration, an IPSec or SSL tunnel is established between the agent and the gateway and a default route is inserted to direct all traffic through the tunnel for the purpose of policy control and threat scanning.
5. The agent submits the client system posture, the HIP.

GlobalProtect allows the use of HIP profiles in security profiles. A HIP object specifies a set of criteria for the client system that is treated as a unit when defining HIP profiles. For example, a HIP object might specify full disk encryption or software patching.

HIP profiles can incorporate HIP objects on a match or no-match basis. For example, a HIP profile might include a match for HIP objects that specify that the client system includes both full disk encryption and certain software patches to be installed.

The portal stores client configurations and maintains internal and external gateways lists. It also lists the CA that is used by gateways and agents for mutual authentication.

Each gateway can operate in tunnel mode (external gateways) or non-tunnel mode (internal gateways). Gateways in non-tunnel mode receive only the HIP from clients. All communication is between agents and gateways; there is no inter-gateway communication.

Policies are enforced on gateways based on user or HIP information, and HIP objects and profile matches are logged in the gateway's HIP database. This information is displayed in the Application Control Center (ACC), logs, and custom reports.

## GlobalProtect Authentication

Connectivity between all parts of the GlobalProtect infrastructure is authenticated using SSL certificates. The portal can act as a certificate authority (CA) for the system (using a self-signed or imported subordinate issuing CA certificate within the portal), or customers can generate certificates using their own CAs. It is recommended (but not required) that the portal, gateways, and the GlobalProtect software agents use certificates signed by the same CA. Prior to transferring any information, the agent verifies the portal server certificate. If the certificate presented by the portal is not signed by a Trusted CA, the agent shows a warning dialog and waits for the user response to continue or cancel.

As part of the configuration bundle that is sent to the client, the portal includes the public certificate of the CA and the needed client certificate and key. The client certificate is used by GlobalProtect gateways to authenticate and identify the client.

If an internal CA is used, the certificate is auto-generated and does not require user interaction. The portal can export the necessary server certificate and key for the gateways. If an external CA is used, support is provided to import the CA certificate along with a server certificate and key for the portal and gateways and a client certificate and key for the clients.

For more information on authentication, refer to "Authentication Profiles" on page 62 and "Authentication Sequence" on page 67.

# Setting Up GlobalProtect

Setting up GlobalProtect on the firewall involves the following tasks:

1. Define HIP objects, as described in "Setting Up HIP Objects" on page 337.
2. Create HIP profiles, as described in "Setting Up HIP Profiles" on page 340.
3. Set up the portal, as described in "Setting Up the GlobalProtect Portal" on page 341.
4. Set up gateways, as described in "Setting Up the GlobalProtect Gateways" on page 347.
5. Define security policies that include HIP profiles, as described in "Defining Security Policies" on page 187.
6. Distribute the GlobalProtect agent, as described in "Setting Up the GlobalProtect Agent" on page 353.
7. Monitor client activity, as described in "Viewing the Logs" on page 264.

### Setting Up HIP Objects

▶ ***Objects > GlobalProtect > HIP Objects***

Use this page to define settings for use in HIP profiles for GlobalProtect. Each HIP object defines a set of criteria for the client system that is treated as a unit when defining HIP profiles.

**Table 137. HIP Object Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Name | Enter a name for the HIP object (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Shared | Select the check box to make the object available to all virtual systems. |
| Description | Enter an optional description. |
| Host Info | Select the check box to specify host information. |
| Domain | To match domains, choose an operator from the drop-down list and enter a string to match. |
| OS | Use the drop-down lists to specify the operating system (OS) for the GlobalProtect agent. |
| Client Versions | To match OS versions on the client, choose an operator from the drop-down list and enter a string to match. |

**Table 137. HIP Object Settings (Continued)**

| Field | Description |
|---|---|
| **Patch Management Tab** | |
| Patch Management | Select the check box to include software patch management in the HIP. When the check box is selected, the settings are activated. You can then define patch management vendors, such as Microsoft, and only allow clients that have certain patch levels installed. |
| Criteria | Specify the following settings on this subtab:<br>• **Is Enabled**—Choose whether the settings on this tab are enabled (yes) or disabled (no), or not available.<br>• **Is Installed**—Select the check box to verify if the client has the defined software patch installed.<br>• **Severity**—Choose the level of importance for missing patches.<br>• **Check**—Choose how the system should check for patches.<br>• **Patches**—Click **Add** and enter patch file names. |
| Vendor | Click **Add** to specify patch management products. Choose a vendor from the drop-down list, and then click **Add** to choose a specific product. Click **OK** to save the settings and return to the **Patch Management** tab. |
| **Firewall Tab** | |
| Firewall | Select the check box to activate this tab and then specify the following settings:<br>• **Is Enabled**—Choose whether the settings on this tab are enabled (yes) or disabled (no), or not available.<br>• **Is Installed**—Select the check box to verify if the client has the defined firewall software installed.<br>• **Vendor and Product**—Click **Add** to specify specific firewalls. Choose a vendor from the drop-down list, and then click **Add** to choose a specific firewall version. Click **OK** to save the settings and return to the **Firewall** tab.<br>• **Exclude Vendor**—Select the check box if you want to exclude rather than include the specified vendors and products. |

**Table 137. HIP Object Settings (Continued)**

| Field | Description |
|---|---|
| **Antivirus Tab** | |
| Antivirus | Select the check box to activate this tab and then specify the following settings:<br>• **Real Time Protection**—Choose whether to require real-time protection (if available).<br>• **Is Installed**—Select the check box to verify if the client has the defined antivirus software installed.<br>• **Virus Definition Version**—Choose from the drop-down list. If you choose **Within** or **Not Within**, specify the number of **Days** or **Hours** to match.<br>• **Product Version**—Choose an operator from the drop-down list and specify a matching string.<br>• **Last Scan Time**—Choose from the drop-down list. If you choose **Within** or **Not Within**, specify the number of **Days** or **Hours** to match.<br>• **Vendor and Product**—Click **Add** to specify antivirus products. Choose a vendor from the drop-down list, and then click **Add** to choose a specific product. Click **OK** to save the settings and return to the **Antivirus** tab.<br>• **Exclude Vendor**—Select the check box if you want to exclude rather than include the specified vendors and products. |
| **Anti-spyware Tab** | |
| Anti-spyware | Select the check box to activate this tab and then specify the following settings:<br>• **Real Time Protection**—Choose whether to require real-time protection (if available).<br>• **Is Installed**—Select the check box to verify if the client has the defined anti-spyware software installed.<br>• **Virus Definition Version**—Choose from the drop-down list. If you choose **Within** or **Not Within**, specify the number of **Days** or **Hours** to match.<br>• **Product Version**—Choose an operator from the drop-down list and specify a matching string.<br>• **Last Scan Time**—Choose from the drop-down list. If you choose **Within** or **Not Within**, specify the number of **Days** or **Hours** to match.<br>• **Vendor and Product**—Click **Add** to specify anti-spyware products. Choose a vendor from the drop-down list, and then click **Add** to choose a specific product. Click **OK** to save the settings and return to the **Anti-spyware** tab.<br>• **Exclude Vendor**—Select the check box if you want to exclude rather than include the specified vendors and products. |

**Table 137. HIP Object Settings (Continued)**

| Field | Description |
|---|---|
| **Disk Backup Tab** | |
| Disk Backup | Select the check box to activate this tab and then specify the following settings:<br>• **Is Installed**—Select the check box to verify if the client has the defined backup software installed.<br>• **Last Backup Time**—Choose from the drop-down list. If you choose **Within** or **Not Within**, specify the number of **Days** or **Hours** to match.<br>• **Vendor and Product**—Click **Add** to specify disk backup products. Choose a vendor from the drop-down list, and then click **Add** to choose a specific product. Click **OK** to save the settings and return to the **Disk Backup** tab.<br>• **Exclude Vendor**—Select the check box if you want to exclude rather than include the specified vendors and products. |
| **Disk Encryption Tab** | |
| Disk Encryption | Select the check box to activate this tab and then specify the following settings: |
| Criteria | Specify the following settings on this subtab:<br>• **Is Installed**—Select the check box to verify if the client has the defined disk encryption software installed.<br>• **Encrypted Locations**—Click **Add** to specify the drive or path that refers to an encrypted data store:<br>– **Encrypted Locations**—Choose the location from the drop-down list.<br>– **State**—Specify the state of the encrypted location by choosing an operator and value from the drop-down list.<br>Click **OK** to save the settings and return to the **Disk Encryption** tab. |
| Vendor | Click **Add** to specify specific disk encryption products. Choose a vendor from the drop-down list, and then click **Add** to choose a specific product. Click **OK** to save the settings and return to the **Disk Encryption** tab. |
| **Custom Checks Tab** | |
| Process List | Click **Add** to specify the list of processes to be checked on the client system to see if they are running. For example, to determine whether a software application is running, add the name of the executable file to the process list. |
| Registry Key | Click **Add** to specify that a particular registry key is present or has a specified value. |
| Plist | Plists are preferences files on MacOS. Define the path to a specific plist file. You can also include preference keys and values to check for within the file. |

## Setting Up HIP Profiles

▶ *Objects > GlobalProtect > HIP Profiles*

After defining HIP objects (refer to "Setting Up HIP Objects" on page 337), use this page to create HIP profiles for GlobalProtect. When defining HIP profiles, you specify match criteria that are built from the previously-defined HIP objects.

**Table 138. HIP Profile Settings**

| Field | Description |
|---|---|
| Name | Enter a name for the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter an optional description. |
| Shared | Select the check box to make the profile available to all virtual systems. |
| Match | Define one or more HIP objects that you want to check for on the client. Enter the HIP objects to include, or click **Add Match Criteria** to create objects. When including multiple HIP objects, you can use AND, OR and NOT operators to create a Boolean expression. Using this method, you can establish complex HIP profiles, for example, to test if your clients have antivirus installed AND disk encryption installed and enabled. |

## Setting Up the GlobalProtect Portal

▶ *Network > GlobalProtect > Portals*

Use this page to configure portals for GlobalProtect.

**Table 139. GlobalProtect Portal Settings**

| Field | Description |
|---|---|
| **Portal Configuration tab** | |
| Name | Enter a name for the portal (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Location | Select the virtual system, if the multiple virtual systems option is enabled. |
| **Network Settings** | |
| Interface | Select the firewall interface that will be used as the ingress for remote clients/firewalls. |
| IP Address | Specify the IP address on which GlobalProtect portal web service will be running. |
| Server Certificate | Select the SSL certificate for the GlobalProtect portal. |
| **Authentication** | |
| Authentication Profile | Choose an authentication profile to authenticate access to the portal. Refer to "Authentication Profiles" on page 62. |
| Client Certificate | Select the certificate the client will use to connect to the gateways. |
| Certificate Profile | Select the certificate profile that is used to authenticate smartcard users on the portal. |
| **Appearance** | |
| Custom Login Page | Choose an optional custom login page for user access to the portal. |
| Custom Help Page | Choose an optional custom help page to assist the user with access to the portal. |

**Table 139. GlobalProtect Portal Settings (Continued)**

| Field | Description |
|---|---|
| **Client Configuration tab** | |

**Table 139. GlobalProtect Portal Settings (Continued)**

| Field | Description |
|---|---|
| General subtab settings | Click **Add** to display the subtabs, and specify the following settings on the **General** subtab:<br>• **Name**—Enter a name to identify this client configuration.<br>• **Use single sign-on**—Select the check box to have GlobalProtect use the users' Windows login credentials to transparently connect and authenticate to the GlobalProtect portal and gateways. No username and password is needed in the GlobalProtect agent configuration.<br>• **Config Refresh Interval (hours)**—Specify the interval in days to refresh the GlobalProtect agent configuration (default 24 hours, range 1-168days).<br>• **Connect Method:**<br>– **on-demand**—Select this option to allow users to establish a connection on demand. With this option, the user must explicitly initiate the connection. This function is primarily used for remote access connections.<br>– **user-logon**—When this option is set, the GlobalProtect agent will automatically establish a connection after users log in to their computers. If you select **Use single sign-on**, the username and password used to log in to Windows is captured by the GlobalProtect agent and used to authenticate.<br>– **pre-logon**—Select this option to preserve pre-logon and post-logon services provided by a corporate infrastructure regardless of where the user machine is located. GlobalProtect will establish a connection prior to user login to the computer. By doing this, a company can create a "logical network" that maintains the security and management features normally achieved by a physical network. Tunnel selection and establishment happens pre-logon based on machine certificates that need to be pre-deployed on client systems outside of GlobalProtect.<br>Examples of some of the services that can be maintained include: Active Directory group policy enforcement, maintaining drive mapping to server resources, and the ability to receive central software deployment downloads while remote. One specific example of how the pre-logon feature works is if a remote user forgets his/her password, since GlobalProtect would connect and use the cached credentials and establish a VPN before the login prompt even appears, a domain administrator could reset the user's password as if they were logged in directly to a domain controller on the physical network. You could also push new Active Directory Group Policies<br><br>With the pre-logon option enabled, you can also use this feature in from the Users tab when configuring Security Policies or in the Source tab in Decryption policies. For example, you may want to set a security policy to only allow access to AD servers and DNS servers on your network for machines that are in the pre-logon state. Note that each user is identified with the name **pre-logon** and you would need to check their machine name in HIP reports to determine specific users.<br>– **Third Party VPN**—Click **Add** to add a list of third party remote access VPN clients that might be present on the system. If configured, GlobalProtect will ignore those clients and their route settings to ensure that it does not interfere or conflict with them. |

**Table 139. GlobalProtect Portal Settings (Continued)**

| Field | Description |
|---|---|
| General subtab (Continued) | – **Internal Host Detection**—With this option, GlobalProtect tries to resolve the configured hostname to the configured IP address. If this fails, GlobalProtect assumes the computer to be outside of the corporate network and will establish a tunnel with any of the available external gateways configured in the **Gateways** tab. Select the check box to enable internal host detection using DNS lookup. Specify the following:<br>– **IP Address**—Enter an internal IP address for the internal host detection.<br>– **Hostname**—Enter the hostname that resolves to the above IP address within the internal network. |
| User/User Group subtab settings | Specify the user or user group to which the particular client configuration is applied. |
| Gateways subtab | Specify gateway settings:<br>• **Cutoff Time**—Specify the timeout (seconds) after which the agent will dismiss a gateway response. The agent dismisses gateway responses after either the configured cutoff time or the socket timeout is reached. If 0 is specified, the cutoff time is ignored by the agent.<br>• **Internal Gateways**—Specify the internal firewalls that the agent will authenticate and provide HIP reports to.<br>• **External Gateways**—Specify the list of firewalls the agent will try to establish a tunnel with when not on the corporate network. The agent will contact all of the gateways and establish a tunnel with the firewall that provides the fastest response and the lowest priority value. Select the **Manual** check box if you want to allow users to manually select a different gateway while their VPN is connected. The GlobalProtect agent will have the option to connect to any external gateway that has the manual selection. When connecting to the new gateway, the existing tunnel will be disconnected and a new tunnel will be established. The manual gateways can also have different authentications mechanism than the primary gateway. If the client system is restarted, or if a rediscovery is performed, the GlobalProtect agent will connect to the primary gateway. This feature is useful if you have a group of users who need to temporarily connect to a specific gateway to access a secure segment of your network.<br><br>**Note:** You can only connect to gateways defined in the portal, you cannot connect to different portals and gateways. |

**Table 139. GlobalProtect Portal Settings (Continued)**

| Field | Description |
|---|---|
| Agent subtab | Specify the following settings:<br>• **Enable advanced view**—Deselect this check box to restrict the user interface on the client side to the basic minimum view. By default, the advanced view setting is enabled on all GlobalProtect agents.<br>• **User can save password**—Select the check box to allow users to save their passwords.<br>• **Agent User Override**—Select an override option:<br>– **disabled**—User override is disabled.<br>– **with-comment**—The user is prompted to enter a comment when disabling the GlobalProtect agent.<br>– **with-passcode**—The user must provide the passcode to use the GlobalProtect agent override.<br>– **with-ticket**—This option enables a challenge-response mechanism to authorize disabling GlobalProtect on the client side. When this option is selected, the user is prompted with a challenge when disabling GlobalProtect. The challenge is then communicated to the firewall administrator out-of-band, and the administrator can validate the challenge through the firewall management interface. The firewall produces a response that is read back to the user who can then disable GlobalProtect by entering the response in GlobalProtect.<br>• **Agent User Override Timeout**—Specify the maximum wait time (seconds) before data collection times out.<br>• **Max Agent User Overrides**—Specify the maximum number of times a user can disable GlobalProtect before a successful connection to a firewall is required.<br>• **Display welcome page**—Select the check box to allow display of a welcome page for the portal.<br>• **Welcome Page**—Choose the factory default welcome page, or click **Import** to import another page. If you choose **None** and select the **Display Welcome Page** option, a blank page is displayed.<br>• **Enable rediscover network option**—Select this check box to allow the user to manually trigger network rediscovery.<br>• **Enable resubmit host information option**—Select this check box to allow the user to manually trigger resubmission of the latest HIP.<br>• **Client Upgrade**—Specify whether to prompt the client to update after configuration changes (**prompt**) or to perform the upgrade silently without a user prompt (**transparent**). |

**Table 139. GlobalProtect Portal Settings (Continued)**

| Field | Description |
|---|---|
| Data Collection Subtab | Specify the following settings on this subtab:<br>• **Max Wait Time**—Specify the maximum wait time (seconds) before data collection times out.<br>• **Exclude Categories**—Click **Add** to specify particular software and client configuration categories to exclude from the data collection. Choose a vendor from the drop down list and click **Add** to choose a specific product. Click **OK** to save settings.<br>• **Custom Checks**—Specify the following information:<br>– **Registry Key**—(Windows) Click **Add** to specify that a particular registry key is present or has a specified value.<br>– **Plist**—(Mac) Click **Add** to specify that a particular plist key is present or has a specified value.<br>– **Process List**—Click **Add** to specify the list of processes to be checked on the end user systems to see if they are running. For example, to determine whether a software application is running, add the name of the executable file to the process list. |
| Trusted Root CA | Specify the Root CA or issuing certificates that the GlobalProtect agent will trust when connecting to a gateway. If a gateway presents a certificate to the client that hasn't been issued by one of the listed CAs, the client will reject the handshake and terminate the connection.<br>Click **Add** to specify a root CA certificate. |
| Agent User Override Key/Confirm Agent User Override Key | When you create a GlobalProtect Agent override with ticket, which can be used to allow users to disable the Agent, you can use this hash key to validate the client when they try to disable with the ticket. |
| **Satellite Configuration** | For information on using GlobalProtect for deploying large scale VPNs, refer to "GlobalProtect Large Scale VPN Deployment" on page 323. |
| General subtab | Click **Add** to display the subtabs, and specify the following on the **GlobalProtect Satellite > General** subtab:<br>• **Name**—Enter a name to identify the GlobalProtect satellite device profile.<br>• **Configuration Refresh Interval (hours)**—Specify how often satellite devices should check the portal for configuration updates (default 24 hours, range 1-48 hours). |
| Devices subtab | Click **Add** to manually add a satellite device using the device serial number. If you use this option, when the satellite device first connects to receive the authentication certificate and the initial configuration, no user login prompt is required. After the satellite device authenticates, the **Name** (host name) will be added automatically to the Portal. |
| Enrollment User/User Group subtab | Click **Add** to add a User or User Group that will be used to authenticate the satellite device when it first connects to the portal. Select **Any** to allow authentication by any user or user group. When the satellite device first connects to the portal, the administrator will be prompted for a login/password. |

**Table 139. GlobalProtect Portal Settings (Continued)**

| Field | Description |
|---|---|
| Gateways subtab | Click **Add** to enter the IP address or hostname of the gateway(s) that will be part of the VPN network. Each satellite device that is added to the portal will create an IPSec tunnel to each gateway. Routing information is also shared between the satellites and gateways.<br>You can also define a routing priority to gateways. When a satellite has VPN tunnels to multiple gateways, duplicate routes may exist. By setting a priority, the gateway with the highest priority will be used first. If the gateway fails, the next highest priority will be used. The priority range is from 1-25, since 25 is the limit on the number of gateways that a satellite can connect to.<br>The satellite will also share its network and routing information with the gateways if the **Publish all static and connect routers to Gateway** (configured on the satellite in Network > IPSec tunnels > GlobalProtect Satellite > Advanced tab) option is selected. Refer to "GlobalProtect Satellite" on page 317.<br>When adding a gateway, in the portal configuration, you create a new key and certificate using the same issuing CA used to create satellite and other gateway certificates. You export the certificate as a PKCS12 file and import it into the gateway together with the issuing CA certificate. |
| Trusted Root CA | Click **Add** and select an existing Trusted Root CA (Certificate Authority), or generate a new Trusted Root CA that will be used to issue client certificates to all satellite and gateway devices for the GlobalProtect satellite profile. |
| Issuing Certificate | Defines the certificate that the portal will use to issue new IPSec client certificates to satellite and gateway devices. |
| Validity Period (days) | Specify the issued GlobalProtect satellite certificate lifetime (default 7 days, range 7-365 days). |
| Certificate Renewal Period (days) | Specify the GlobalProtect satellite certificate renewal period (default 3 days, range 3-30 days). This will determines how often certificates should be renewed. |
| OCSP Responder | Select the OCSP responder that will be used by the portal to determine the revocation status of the gateway and satellite devices. |

## Setting Up the GlobalProtect Gateways

▶ *Network > GlobalProtect > Gateways*

Use this page to configure gateways for GlobalProtect. The gateway can be used to provide VPN connections for client PCs or GlobalProtect satellite devices.

**Table 140. GlobalProtect Gateway Settings**

| Field | Description |
|---|---|
| **General tab** | |
| Name | Enter a name for the gateway (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Location | Select the virtual system, if the multiple virtual systems option is enabled. |
| Network Settings | |

**Table 140. GlobalProtect Gateway Settings (Continued)**

| Field | Description |
|---|---|
| Interface | Select the firewall interface that will be used as the ingress for remote clients/firewalls. |
| IP Address | Specify the IP address on which GlobalProtect portal web service will be running. |
| Server Certificate | Choose the server certificate for the gateway. |
| Authentication | |
| Authentication Profile | Choose an authentication profile or sequence to authenticate access to the portal. Refer to "Authentication Profiles" on page 62. |
| Certificate Profile | Choose the certificate profile for client authentication. |
| **Client Configuration tab** | |
| Tunnel Settings subtab | |
| Tunnel Mode | Select the check box to enable tunnel mode and specify the following settings:<br>• **Tunnel Interface**—Choose the tunnel interface for access to the gateway.<br>• **Max Users**—Specify the maximum number of users that can access the gateway at the same time. If the maximum number of users is reached, subsequent users are denied access with an error message indicating that the maximum number of users has been reached.<br>• **Enable IPSec**—Select the check box to enable IPSec mode for client traffic, making IPSec the primary and SSL-VPN the fall back method.<br>• **Enable X-Auth Support**—Select the check box to enable Extended Authentication (X-Auth) support in the GlobalProtect gateway when IPSec is enabled. With X-Auth support, third party IPSec VPN clients that support X-Auth (such as the IPSec VPN client on Apple iOS and Android devices) can establish a VPN tunnel with the GlobalProtect gateway. The X-Auth option simply provides remote access from the VPN client to a specific GlobalProtect gateway and does not provide the full control and HIP features that are part of the GlobalProtect agent installed on PCs. For details on configuring a GlobalProtect gateway to allow IPSec X-Auth clients, refer to the tech notes "GlobalProtect Config for Apple iOS VPN" and "GlobalProtect Config for Android VPN" in the Tech Note area at *https://live.paloaltonetworks.com/community/documentation.*<br>– If the group name and group password are specified, the first authentication phase requires both parties to use this credential to authenticate. The second phase requires a valid user name and password, which is verified through the authentication profile configured in the Authentication section.<br>– If no group name and group password are defined, the first authentication phase is based on a valid certificate presented by the third party VPN client. This certificate is then validated through the certificate profile configured in the authentication section. |

**Table 140. GlobalProtect Gateway Settings (Continued)**

| Field | Description |
|---|---|
| Timeout Configuration | Specify the following timeout settings:<br>• **Login Lifetime**—Specify the number of days, hours, or minutes allowed for a single gateway login session.<br>• **Inactivity Logout**—Specify the number of days, hours, or minutes after which an inactive session is automatically logged out. |
| **Network Settings subtab** | |
| Inheritance Source | Select a source to propagate DNS server and other settings from the selected DHCP client or PPPoE client interface into the GlobalProtect agents' configuration. With this setting all client network configuration, such as DNS servers and WINS servers, are inherited from the configuration of the interface selected in the Inheritance Source. |
| Primary DNS<br>Secondary DNS | Enter the IP addresses of the primary and secondary servers that provide DNS to the clients. |
| Primary WINS<br>Secondary WINS | Enter the IP addresses of the primary and secondary servers that provide Windows Internet Naming Service (WINS) to the clients. |
| Check inheritance source status | Click the link to see the server settings that are currently assigned to the client interfaces. |
| | Click **Add** to enter a suffix that the client should use locally when an unqualified hostname is entered that it cannot resolve. You can enter multiple suffixes by separating them with commas. |
| Inherit DNS Suffixes | Select this check box to inherit the DNS suffixes from the inheritance source. |
| IP Pool | Click **Add** to specify IP pool settings.<br>Use this section to create a range of IP addresses to assign to remote users. When the tunnel is established, an interface is created on the remote user's computer with an address in this range.<br>***Note:*** *The IP pool must be large enough to support all concurrent connections. IP address assignment is dynamic and not retained after the user disconnects. Configuring multiple ranges from different subnets will allow the system to offer clients an IP address that does not conflict with other interfaces on the client.*<br>The servers/routers in the networks must route the traffic for this IP pool to the firewall.<br>For example, for the 192.168.0.0/16 network, a remote user may be assigned the address 192.168.0.10. |
| Access Route | Click **Add** to specify access route options.<br>Use this section to add routes that will be pushed to the remote user's computer and therefore determine what the user's computer will send through the VPN connection.<br>For example, you can set up split tunneling to allow remote users to access the Internet without going through the VPN tunnel.<br>If no route is added, then every request is routed through the tunnel (no split tunneling). In this case, each Internet request passes through the firewall and then out to the network. This method can prevent the possibility of an external party accessing the user's computer and then gaining access to the internal network (with the user's computer acting as bridge). |

**Table 140. GlobalProtect Gateway Settings (Continued)**

| Field | Description |
|---|---|
| **HIP Notification subtab** | |
| HIP Notification | Click **Add** to specify notification options. Select **Enable** to enable the option to match or not match a message.<br>Choose a notification option from the **Show Notification As** section and choose the radio button for a **System Tray Balloon** or **Pop Up Message**, and then specify a message to match or not match. Use these settings to notify the end user about the state of the machine, for example, to provide a warning message.<br>***Note:*** *The HIP notification pages can be formatted in rich HTML, which can include links to external web sites and resource. Use the link icon* *in the rich text settings toolbar to add links.* |
| **Satellite Configuration tab** | |
| **Tunnel Settings subtab** | |
| Tunnel Configuration | Select the **Tunnel Configuration** check box and select an existing tunnel interface, or click New Tunnel Interface. For information on creating a tunnel interface, refer to "Configuring Tunnel Interfaces" on page 147.<br>**Replay attack detection**—Protect against replay attacks.<br>**Copy TOS**—Copy the (Type of Service) TOS header from the inner IP header to the outer IP header of the encapsulated packets in order to preserver the original TOS information.<br>**Configuration refresh interval (hours)**—Specify how often satellite devices should check the portal for configuration updates (default 24 hours, range 148 hours). |
| Tunnel Monitoring | Select the **Tunnel Monitoring** check box to enable monitoring of the VPN tunnels between the satellite device and gateways. This is recommended to allow failover if one of the gateways cannot communicate with the satellite.<br>**Destination IP**—Specify an IP address on the other side of the tunnel that the tunnel monitor will use to determine if the tunnel is working properly.<br>**Tunnel Monitor Profile**—Select the default monitor profile, or create a new profile. A monitor profile is used to perform automatic failover to another tunnel. Refer to "Defining Monitor Profiles" on page 177. |
| Crypto Profiles | Select an **IPSec Crypto Profile**, or create a new profile. This will determine the protocols and algorithms for identification, authentication, and encryption for the VPN tunnels. Refer to "Defining IPSec Crypto Profiles" on page 319. |
| **Network Settings subtab** | |
| Inheritance Source | Select a source to propagate DNS server and other settings from the selected DHCP client or PPPoE client interface into the GlobalProtect satellite configuration. With this setting all network configuration, such as DNS servers, are inherited from the configuration of the interface selected in the Inheritance Source. |

**Table 140. GlobalProtect Gateway Settings (Continued)**

| Field | Description |
|---|---|
| Primary DNS<br>Secondary DNS | Enter the IP addresses of the primary and secondary servers that provide DNS to the satellites. |
| DNS Suffix | Click **Add** to enter a suffix that the satellite should use locally when an unqualified hostname is entered that it cannot resolve. You can enter multiple suffixes by separating them with commas. |
| Inherit DNS Suffix | Select this check box to send the DNS suffix to the satellite devices to use locally when an unqualified hostname is entered that it cannot resolve. |
| IP Pool | Click **Add** to specify IP pool settings.<br>Use this section to create a range of IP addresses to assign to satellite devices. When the tunnel is established, an interface is created on the satellite device with an address in this range.<br>***Note:*** *The IP pool must be large enough to support all concurrent connections. IP address assignment is dynamic and not retained after the satellite disconnects. Configuring multiple ranges from different subnets will allow the system to offer satellites an IP address that does not conflict with other interfaces on the device.*<br>The servers/routers in the networks must route the traffic for this IP pool to the firewall.<br>For example, for the 192.168.0.0/16 network, a satellite may be assigned the address 192.168.0.10. |
| Access Route | Click **Add** to specify access route options.<br>Use this section to add routes that will be pushed to the satellites and therefore determining what routes the satellite will use to send traffic through the VPN connection.<br>For example, you can set up split tunneling to allow satellites to access the Internet without going through the VPN tunnel.<br>If no route is added, then every request is routed through the tunnel (no split tunneling).<br>Click **Add** to enter a route. |
| Route Filter subtab | Click the **Accept published routes** check box and enter the satellite subnets that will be added to the gateway's route table when routing information is shared. Satellite subnets that are not part of the list will be filtered out. |

# Setting Up and Activating the GlobalProtect Agent

▶ *Device > GlobalProtect Client*

The **GlobalProtect Client** page lists the available GlobalProtect releases. When the client connects, the system checks the version and installs the currently activated version of the agent software if it is different from the version that is on the client.

***Note:*** *For initial download and installation of the GlobalProtect agent, the user on the client system must be logged in with administrator rights. For subsequent upgrades, administrator rights are not required.*

To download and activate the GlobalProtect agent:

1. Click the **Download** link for the desired release. The download starts and a pop-up window opens to display the progress of the download. When the download is complete, click **Close**.
2. To activate a downloaded release, click the **Activate** link for the release. If an existing version of the client software has already been downloaded and activated, a pop-up message is displayed to indicate that the new version will be downloaded the next time that the clients connect.
3. To activate the client that was previously installed by way of the **Upload** button, click the **Activate From File** button. A pop-up window opens. Select the file from the drop-down list and click **OK**.
4. To remove a downloaded release of the client software from the firewall, click the **Remove** icon in the rightmost column.

# Setting Up the GlobalProtect Agent

The GlobalProtect agent (PanGP Agent) is an application that is installed on the client system (typically a laptop) to support GlobalProtect connections with portals and gateways and is supported by the GlobalProtect service (PanGP Service).

***Note:*** *Make sure that you choose the correct installation option for your host operating system (32-bit or 64-bit). If installing on a 64-bit host, use 64-bit browser/Java combo for the initial installation.*

To install the agent, open the installer file and follow the on-screen instructions.

To configure the agent:

1. Choose **Start > All Programs > Palo Alto Networks > GlobalProtect > GlobalProtect**.

   The client interface opens to show the **Settings** tab.

**Figure 45. GlobalProtect - Settings Tab**

2. Specify the username and password to use for GlobalProtect authentication, and optionally select the **Remember Me** check box.
3. Enter the IP address of the firewall that serves as the GlobalProtect Portal.
4. Click **Apply**.

### Using the GlobalProtect Agent

The tabs in the GlobalProtect agent contain useful information about status and settings, and provide information to assist in troubleshooting connection issues.

- **Status tab**—Displays current connection status and lists any warnings or errors.
- **Details tab**—Displays information about the current connection, including portal IP addresses and protocol, and presents byte and packet statistics about the network connection.
- **Host State tab**—Displays the information stored in the HIP. Click a category on the left side of the window to display the configured information for that category on the right side of the window.

- **Troubleshooting tab**—Displays information to assist in troubleshooting.
  - **Network Configurations**—Displays the current client system configuration.
  - **Routing Table**—Displays information on how the GlobalProtect connection is currently routed.
  - **Sockets**—Displays socket information for the current active connections.
  - **Logs**—Allows you to display logs for the GlobalProtect agent (PanGP Agent) and service (PanGP Service). Choose the log type and debugging level. Click **Start** to begin logging and **Stop** to terminate logging.

Chapter 10

# Configuring Quality of Service

This chapter describes how to configure quality of service (QoS) on the firewall:

- "Firewall Support for QoS" in the next section
- "Defining QoS Profiles" on page 358
- "Defining QoS Policies" on page 359
- "Displaying QoS Statistics" on page 362

## Firewall Support for QoS

The firewall supports fine grained QoS settings for clear text and tunneled traffic upon egress from the firewall. (Ingress QoS processing is not supported.) QoS profiles are attached to physical interfaces to specify how traffic classes map to bandwidth (guaranteed, maximum) and priority. QoS policy is then used to map specific sessions to QoS classes. QoS classification is supported with all interface types except Aggregate Ethernet.

The firewall supports the following QoS settings:

- Use the Network > QoS page to configure QoS settings for firewall interfaces and specify criteria for the clear text and tunneled traffic that leaves the firewall through those interfaces. Refer to "Configuring QoS for Firewall Interfaces" on page 356.
- For each interface, you can define QoS profiles that determine how the QoS traffic classes are treated. You can set overall limits on bandwidth regardless of class and also set limits for individual classes. You can also assign priorities to different classes. Priorities determine how traffic is treated when contention occurs. Refer to "Defining QoS Profiles" on page 358.
- Use the Policies > QoS page to configure policies to activate the QoS restrictions. Refer to "Defining QoS Policies" on page 359.

### Important items to consider when configuring firewall support for QoS

- When setting up the QoS profile, the guaranteed and maximum egress settings defined for the classes must not exceed the guaranteed and maximum egress settings defined for the profile itself.
- Traffic that does not match QoS policy will be assigned a default class of 4. Be sure to assign a maximum guaranteed bandwidth and priority with this in mind.
- Each firewall model supports a maximum number of ports that can be configured with QoS. Refer to the spec sheet for your firewall model at *http://www.paloaltonetworks.com*.

# Configuring QoS for Firewall Interfaces

▶ *Network > QoS*

Use the **QoS** page to configure bandwidth limits for firewall interfaces.

**Table 141. QoS Settings**

| Field | Description |
|---|---|
| **Physical Interface** | |
| Interface Name | Select the firewall interface. |
| Egress Max (Mbps) | Enter the limit on traffic leaving the firewall through this interface (Mbps). |
| Turn on QoS feature on this interface | Select the check box to enable QoS features. |
| Default Profile:<br>Clear Text<br>Tunnel Interface | Select the default QoS profiles for clear text and for tunneled traffic. You must specify a default profile for each. For clear text traffic, the default profile applies to all clear text traffic as an aggregate. For tunneled traffic, the default profile is applied individually to each tunnel that does not have a specific profile assignment in the detailed configuration section. For instructions on defining QoS profiles, refer to "Defining QoS Profiles" on page 358. |
| **Clear Text Traffic and Tunneled Traffic** | Specify the following settings on the **Tunneled** and **Clear Text Traffic** tabs. These values apply unless they are overridden by setting in the Detail Configuration area, as described later in this table. |
| Egress Guaranteed (Mbps) | Enter the bandwidth that is guaranteed for tunneled traffic from this interface. |
| Egress Max (Mbps) | Enter the limit on traffic leaving the firewall through this interface (Mbps). |

**Table 141. QoS Settings (Continued)**

| Field | Description |
|---|---|
| Add | Click **Add** from the **Clear Text Traffic** or **Tunneled Traffic** tabs to define additional granularity to the treatment of clear text traffic or to override the default profile assignment for specific tunnels. If this section is left blank, the values specified in Group Configuration are used.<br>For example, assume a configuration with two sites, one of which has a 45 Mbps connection and the other a T1 connection to the firewall. You can apply restrictive QoS settings to the T1 site so that the connection is not overloaded while also allowing more flexible settings for the site with the 45 Mbps connection.<br>To add granularity for clear text traffic, click the **Clear Text** tab, click **Add**, and then click individual entries to configure the following settings:<br>• **Name**—Enter a name to identify these settings.<br>• **QoS Profile**—Select the QoS profile to apply to the specified interface and subnet. For instructions on defining QoS profiles, refer to "Defining QoS Profiles" on page 358.<br>***Note:*** *The QoS rules for clear text are applied in the specified order. Refer to the following tech note for more details on QoS https://live.paloaltonetworks.com/docs/DOC-3439.*<br>• **Source Interface**—Select the firewall interface.<br>• **Source Subnet**—Select a subnet to restrict the settings to traffic coming from that source, or keep the default **any** to apply the settings to any traffic from the specified interface.<br>To override the default profile for a specific tunnel, click the **Tunneled Traffic** tab, click **Add**, and then click individual entries to configure the following settings:<br>• **Tunnel Interface**—Select the tunnel interface on the firewall.<br>• **QoS Profile**—Select the QoS profile to apply to the specified tunnel interface.<br>To remove a clear text or tunneled traffic entry, select the check box for the entry and click **Remove**. |

# Defining QoS Profiles

▶ *Network > Network Profiles > QoS Profiles*

For each interface, you can define QoS profiles that determine how the QoS traffic classes are treated. You can set overall limits on bandwidth regardless of class and also set limits for individual classes. You can also assign priorities to different classes. Priorities determine how traffic is treated in the presence of contention.

***Note:*** *Refer to "Configuring QoS for Firewall Interfaces" on page 356 for information on configuring firewall interfaces for QoS and refer to "Defining QoS Policies" on page 359 to configure the policies that will activate the QoS restrictions.*

**Table 142. QoS Profile Settings**

| Field | Description |
|---|---|
| Profile Name | Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Egress Max | Enter the maximum bandwidth allowed for this profile (Mbps). |
| Egress Guaranteed | Enter the bandwidth that is guaranteed for this profile (Mbps). |
| Classes | Specify how to treat individual QoS classes. You can select one or more classes to configure:<br>• **Class**—If you do not configure a class, you can still include it in a QoS policy. In this case, the traffic is subject to overall QoS limits. Traffic that does not match a QoS policy will be assigned to class 4.<br>• **Priority**—Click and select a priority to assign to this class. These are prioritized in the order listed (highest first):<br>– real-time<br>– high<br>– medium<br>– low<br>• **Egress Max**—Click and enter a value (Mbps) for this class.<br>• **Egress Guaranteed**—Click and enter a value (Mbps) for this class.<br>When contention occurs, traffic that is assigned a lower priority is dropped. Real-time priority uses its own separate queue. |

# Defining QoS Policies

▶ *Policies > QoS*

The QoS policy determines how traffic is classified for treatment when it passes through an interface with QoS enabled. For each rule, you specify one of eight classes. You can also assign a schedule to specify which rule is active. Unclassified traffic is automatically assigned to class 4.

***Note:*** *Refer to "Configuring QoS for Firewall Interfaces" on page 356 for information on configuring firewall interfaces for QoS and refer to "Defining QoS Profiles" on page 358 for information on configuring classes of service.*

To view just the rules for a specific virtual system, select the system from the **Virtual System** drop-down list and click **Go**. To apply a filter to the list, select from the **Filter Rules** drop-down list. To view just the rules for specific zones, select a zone from the **Source Zone** and/or **Destination Zone** drop-down lists, and click **Filter by Zone**.

***Note:*** *Shared polices pushed from Panorama are shown in green and cannot be edited at the device level.*

To add a new QoS rule, do one of the following:

- Click **Add** at the bottom of the page and configure the rule. A new rule is added to the bottom of the list.
- Select **Clone Rule**, or select a rule by clicking the white space of the rule, and select **Clone** at the bottom of the page (a selected rule has a yellow background). The copied rule is inserted below the selected rule.

**Table 143. QoS Rule Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Name | Enter a name to identify the rule (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter an optional description. |
| Tag | If you need to tag the policy, click **Add** to specify the tag.<br>A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location. |
| **Source Tab** | |

**Table 143. QoS Rule Settings (Continued)**

| Field | Description |
|---|---|
| Source Zone | Select one or more source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151. |
| Source Address | Specify a combination of source IPv4 or IPv6 addresses for which the identified application can be overridden. To select specific addresses, choose **select** from the drop-down list and do any of the following:<br>• Select the check box next to the appropriate addresses and/or address groups in the **Available** column, and click **Add** to add your selections to the **Selected** column.<br>• Enter the first few characters of a name in the **Search** field to list all addresses and address groups that start with those characters. Selecting an item in the list will set the check box in the **Available** column. Repeat this process as often as needed, and then click **Add**.<br>• Enter one or more IP addresses (one per line), with or without a network mask. The general format is:<br><ip_address>/<mask><br>• To remove addresses, select the appropriate check boxes in the **Selected** column and click **Remove**, or select **any** to clear all addresses and address groups.<br>To add new addresses that can be used in this or other policies, click **New Address** (refer to "Defining Applications" on page 233). To define new address groups, refer to "Defining Address Groups" on page 230. |
| Source User | Specify the source users and groups to which the QoS policy will apply. |
| Negate | Select the check box to have the policy apply if the specified information on this tab does NOT match. |
| **Destination Tab** | |
| Destination Zone | Select one or more source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151. |
| Destination Address | Specify a combination of source IPv4 or IPv6 addresses for which the identified application can be overridden. To select specific addresses, choose **select** from the drop-down list and do any of the following:<br>• Select the check box next to the appropriate addresses and/or address groups in the **Available** column, and click **Add** to add your selections to the **Selected** column.<br>• Enter the first few characters of a name in the **Search** field to list all addresses and address groups that start with those characters. Selecting an item in the list will set the check box in the **Available** column. Repeat this process as often as needed, and then click **Add**.<br>• Enter one or more IP addresses (one per line), with or without a network mask. The general format is:<br><ip_address>/<mask><br>• To remove addresses, select the appropriate check boxes in the **Selected** column and click **Remove**, or select **any** to clear all addresses and address groups.<br>To add new addresses that can be used in this or other policies, click **New Address** (refer to "Defining Applications" on page 233). To define new address groups, refer to "Defining Address Groups" on page 230. |

**Table 143. QoS Rule Settings (Continued)**

| Field | Description |
|---|---|
| Negate | Select the check box to have the policy apply if the specified information on this tab does NOT match. |
| **Application Tab** | |
| Application | Select specific applications for the QoS rule. To define new applications, refer to "Defining Applications" on page 233. To define application groups, refer to "Defining Application Groups" on page 238.<br>If an application has multiple functions, you can select the overall application or individual functions. If you select the overall application, all functions are included, and the application definition is automatically updated as future functions are added.<br>If you are using application groups, filters, or container in the QoS rule, you can view details on these objects by holding your mouse over the object in the **Application** column, click the down arrow and select **Value**. This enables you to easily view application members directly from the policy without having to go to the Object tabs. |
| **Service/ URL Category Tab** | |
| Service | Select services to limit to specific TCP and/or UDP port numbers. Choose one of the following from the drop-down list:<br>• **any**—The selected applications are allowed or denied on any protocol or port.<br>• **application-default**—The selected applications are allowed or denied only on their default ports defined by Palo Alto Networks. This option is recommended for allow policies.<br>• **Select**—Click **Add**. Choose an existing service or choose **Service** or **Service Group** to specify a new entry. Refer to "Services" on page 239 and "Service Groups" on page 240. |
| URL Category | Select URL categories for the QoS rule.<br>• Choose **any** to ensure that a session can match this QoS rule regardless of the URL category.<br>• To specify a category, click **Add** and select a specific category (including a custom category) from the drop-down list. You can add multiple categories. Refer to "Custom URL Categories" on page 242 for information on defining custom categories. |
| **Other Settings Tab** | |
| Class | Choose the QoS class to assign to the rule, and click **OK**. Class characteristics are defined in the QoS profile. Refer to "Defining QoS Profiles" on page 358 for information on configuring settings for QoS classes. |
| Schedule | Choose the calendar icon to set a schedule for the QoS policy to apply. |

# Displaying QoS Statistics

▶ *Network > QoS*

The table on the **QoS Policies** page indicates when QoS is enabled, and includes a link to display QoS statistics. An example is shown in the following figure.

**Figure 46. QoS Statistics**

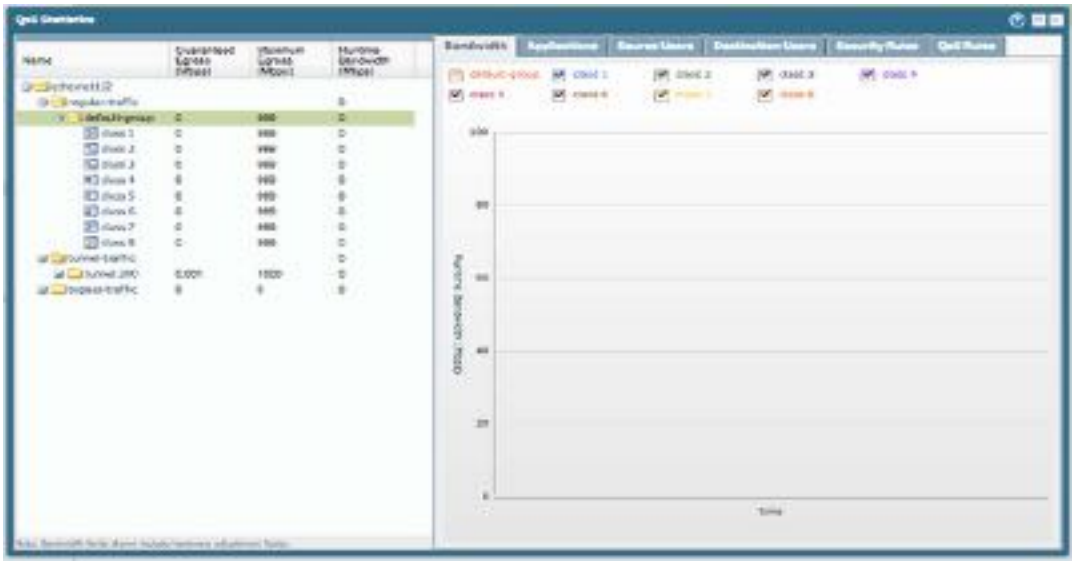

The left panel shows the QoS tree table, and the right panel shows data in the following tabs:

- **QoS Bandwidth**—Shows the real time bandwidth charts for the selected node and classes. The information is updated every two seconds.
- **Session Browser**—Lists the active sessions of the selected node and/or class.
- **Application View**—Lists all active applications for the selected QoS node and/or class.

Chapter 11

# Setting Up a VM-Series Firewall

This chapter describes how to install the VM-Series firewall:


- "Overview" in the next section
- "System Requirements and Limitations" on page 364
- "About Licensing the VM-Series Firewall" on page 365
- "Installing and Licensing the VM-Series Firewall" on page 366
- "Troubleshooting" on page 370


## Overview

The Palo Alto Networks VM-Series firewall is the virtualized form of the Palo Alto Networks firewall. It is positioned for use in a virtualized data center environment and is particularly well suited for private and public cloud deployments. By using virtual machine technology, you can install this solution on any x86 device that is capable of running VMware ESXi, without the need to deploy Palo Alto Networks hardware.

The VM-Series firewall has many characteristics in common with the Palo Alto Networks hardware firewalls, including common features and management interfaces. The main difference is in the deployment methods used in a virtual environment in order to fit your needs. Once the virtual firewall is installed, you will use it and manage it much like you do with the hardware firewalls.

The VM-Series firewall is distributed using the Open Virtualization Format (OVF), which is a standard method of packaging and deploying virtual machines.

***Note:*** This section covers the basic installation steps needed to deploy a VM-Series firewall. For more details and uses cases, refer to the VM-Series firewall Tech Note at *https://live.paloaltonetworks.com/community/documentation*.

# System Requirements and Limitations

This section lists requirements and limitations for the VM-Series firewall.

## Requirements

You can create and deploy multiple instances of the VM-Series firewall on an ESX(i) server. Because each instance of the firewall requires a minimum resource allocation — number of CPUs, memory and disk space— on the ESX(i) server, make sure to conform to the specifications below to ensure optimal performance.

The VM-Series firewall has the following requirements:

- VMware ESX(i) with vSphere 4.1 and 5.0.
- Minimum of two vCPUs per VM-Series firewall. One will be used for the for the management plane and one for the dataplane. You can add up to eight additional vCPUs for the dataplane in the following increments: 2, 4, or 8 vCPUs.
- Minimum of two network interfaces (vmNICs). One will be a dedicated vmNIC for the management interface and one for the data port. You can then add up to eight more vmNICs for data traffic.
  The VM-Series firewall requires that promiscuous mode is set to "accept" on the port group of the virtual switch to which the data ports on the firewall are attached.
- Minimum of 4GB of memory. Any additional memory will be used by the management plane only.
- Minimum of 40GB of virtual disk space. You can add an additional disk of up to 2TB for logging purposes.

## Limitations

The VM-Series firewall functionality is very similar to the Palo Alto Networks hardware firewalls, but with the following limitations:

- Dedicated CPU cores are required.
- Only High Availability (HA) lite is supported (active/passive with no stateful failover).
- High Availability (HA) "Link Monitoring" is only supported on VMware ESXi installations that support DirectPath I/O.
- Up to 10 total ports can be configured; this is a VMware limitation. One port will be used for management traffic and up to 9 can be used for data traffic.
- Only the vmxnet3 driver is supported.
- Virtual systems are not supported.
- vMotion is not supported.
- Jumbo frames are not supported.
- Link Aggregation is not supported.

# About Licensing the VM-Series Firewall

The VM-Series firewall is available in three models — VM-100, VM-200, and VM-300— and each model is licensed for a maximum capacity. Capacity is defined in terms of the number of sessions, rules, security zones, address objects, IPSec VPN tunnels and SSL VPN tunnels that the VM-Series firewall is optimized to handle. When purchasing a license, make sure to purchase the correct model for your network requirements. While each model has a distinct orderable part number, the software package that is used to create the VM-Series firewall is the same across all models.

When you purchase a VM-Series firewall, you will receive a set of auth-codes over email. Typically the email includes a capacity auth-code for the model purchased (VM-100, VM-200, VM300), a software and support auth-code (for example, PAN-SVC-PREM-VM-100 SKU auth-code) that provides access to software/content updates and support. If you purchased additional subscriptions for Threat Prevention, URL Filtering, GlobalProtect, or WildFire, a list of the other auth-codes purchased with the order are included.

If you do not have an existing support account, you must use the capacity auth-code to register and create an account on the Support Portal. After your account is verified and the registration is complete, you will be able to log in and download the software package required to install the VM-Series firewall. If you have an existing support account, you can access the "VM-Series Authentication Code" link on the support software page to manage your VM-Series firewall licenses and download the software.

***Note:*** *If you have an evaluation copy of the VM-Series firewall and would like to convert it to a fully licensed (purchased) copy, clone your VM-Series firewall and use the instructions in the following sections to register and license the purchased copy of your VM-Series firewall.*

# Registering the VM-Series Firewall

You are required to register your capacity auth-code with your support account.

1. Log in to https://support.paloaltonetworks.com with your account credentials.
2. Select **My VM-Series Auth-Codes** link on the home page and click **Add VM-Series Auth-Code.**
3. Enter the capacity auth-code you received by email, and click the **Add VM-Series Auth-Code** button. The page will refresh and you will see the list of auth-codes registered to your support account.
4. Continue with "Installing and Licensing the VM-Series Firewall" on page 366.

# Installing and Licensing the VM-Series Firewall

This section describes the steps needed to install and license the VM-Series firewall.

Installing a VM-Series firewall:

1. Obtain a VM-Series license then download the Open Virtualization Format (OVF) template from *https://support.paloaltonetworks.com* software site.

   The OVF is downloaded as a zip archive that is expanded into three files:

   – OVF extension—The OVF descriptor file that contains all metadata about the package and its contents.

   – MF extension—The OVF manifest file that contains the SHA-1 digests of individual files in the package.

   – VMDK extension—The disk image file that contains the virtualized version of the PAN-OS firewall.

   ***Note:*** *The above files contain the base installation. After the base installation is complete, you will need to download and install the latest version of the VM-Series firewall software from the support site. This will ensure that you have the latest fixes that were implemented since the base OVF file was created.*

2. Before deploying the template, it is helpful to setup virtual standard switch(es) and virtual distributed switch(es) that you will need for the VM-Series firewall. The VM-Series firewall requires that any attached virtual switch has promiscuous mode enabled.

   To configure a virtual standard switch:

   a. Configure a virtual standard switch from the vSphere Client by navigating to **Home > Inventory > Hosts and Clusters**.

   b. Click the **Configuration** tab and under **Hardware** click **Networking.** For each VM-Series firewall attached virtual switch, click on **Properties**.

   c. Highlight the virtual switch and click **Edit**. In the vSwitch properties, click the **Security** tab and set **Promiscuous Mode:** to **Accept** and then click **OK**. This change will propagate to all port groups on the virtual switch.

   To configure a virtual distributed switch:

   a. Navigate to **Home > Inventory > Networking**. Highlight the **Distributed Port Group** you want to edit and select the **Summary** tab.

   b. Click on **Edit Settings** and select **Policies > Security** and set **Promiscuous Mode:** to **Accept** and then click **OK**.

3. Deploy the OVF template:

   a. Log in to vCenter using the vSphere client. You can also go directly to the target ESXi host if needed.

   b. From the vSphere client, select **File > Deploy OVF Template**.

c. Browse to the OVF template that you downloaded in step 1, select the file and then click **Next**. Review the templates details window and then click **Next** again.

d. Name the VM-Series firewall instance and in the **Inventory Location:** window, select a Data Center and folder and click **Next**.

e. Select an ESXi host for the VM-Series firewall and click **Next**.

f. Select the datastore to use for the VM-Series firewall and click **Next**.

g. Leave the default settings for the datastore provisioning and click **Next**. The default is "Thick Provision Lazy Zeroed".

h. Select the networks to use for the two initial vmNICs. The first vmNIC will be used for the management interface and the second vmNIC for the first data port. Make sure that the **Source Networks** is mapped to the correct **Destination Networks**.

i. Review the details window, click the **Power on after deployment** check box and then click **Next**.

You can monitor the **Recent Tasks** list to view the progress of the deployment. When the deployment is complete, click the **Summary** tab to review the current status.

4. To perform the initial configuration of the VM-Series firewall, follow these steps:
   a. The default management IP is 192.168.1.1. To change this, launch the vSphere console, click the **Summary** tab and under **Commands** click **Open Console**. You can also right click the VM and select Open Console.
   b. Log in with the default login user name **admin** password **admin** and then type **configure** to enter configuration mode.
   c. Set the desired IP, netmask, gateway, and DNS IP. Example:
      ```
      set deviceconfig system ip-address 10.1.1.5 netmask
      255.255.255.0 default-gateway 10.1.1.1 dns-setting servers
      primary 10.0.0.245
      ```
   d. Type **commit** to make the changes active.
   e. Test network connectivity to your default gateway and a known server. Example: `ping host 10.1.1.1` or `ping host 10.0.0.245`.
5. Activate the license. Pick one of the following options based on whether or not the VM-Series firewall has direct internet access.
   a. If your firewall has direct internet access:
      i. Navigate to the **Device >Licenses** tab and select the **Activate feature using authentication code** link.
      ii. Enter the capacity auth-code that you registered on the Support website. The firewall will connect to the update server (updates.paloaltonetworks.com), and download the license and reboot automatically.
      iii. Log back in to the web interface and confirm that the **Dashboard** displays a valid serial number. If the term *Unknown* displays, it means the device is not licensed.
      iv. On the **Device-> Licenses** tab, verify that "PA-VM" license is added to the device.

b. In case the ESXi server that your VM-Series firewall is installed on does not have internet access:

   i. Navigate to the **Device > Licenses** tab and click the **Activate Feature using Auth Code** link.

   ii. Click D**ownload Authorization File**, and download the *authorizationfile.txt* on the client machine.

   iii. Copy the *authorizationfile.txt* to a computer that has access to the internet and log in to the support portal. Click **My VM-Series Auth-Codes** link and select the applicable auth-code from the list and click the **Register V**M link.

   iv. On the **Register VM Device** tab upload the authorization file. This will complete the registration process and the serial number of your VM-Series firewall will be attached to your account records.

   v. Navigate to the **My Devices** tab and search for the VM-Series device just registered and click the **PA-VM** link. This will download the VM-Series license key to the client machine.

   vi. Copy the license key to the machine that can access the web interface of the VM-Series firewall and navigate to the **Device -> Licenses** tab.

   vii. Click **Manually Upload License** link and enter the license key. The license will be activated on the device and the device will reboot.

   viii. Log in to the device and confirm that the **Dashboard** displays a valid serial number and that the PA-VM license displays in the **Device > Licenses** tab.

***Note:*** *In order to complete the license installation process, you must reboot your VM-Series firewall instance.*

6. Now that the VM-Series firewall has network connectivity and the base PAN-OS software is installed, you need to upgrade to the latest version of PAN-OS (a support license is required).

***Note:*** *The base VM-Series firewall software is installed with the OVF file. It's recommended that you upgrade to the latest content update and PAN-OS release to ensure you have all of the latest fixes. At each feature release, a new base image will be created as part of the OVF file.*

   a. From the web interface, navigate to **Device > Licenses** and make sure you have the correct VM-Series firewall license and that the license is activated.

   b. To upgrade the VM-Series firewall PAN-OS software, navigate to **Device > Software**.

   c. Click **Refresh** to view the latest software release and also review the **Release Notes** to view a description of the changes in a release and to view the migration path to install the software.

   d. Click **Download** to retrieve the software then click **Install** to install it.

   For details on downloading and installing PAN-OS software, refer to "Upgrading/ Downgrading the PAN-OS Software" on page 50.

# Troubleshooting

Many of the troubleshooting steps for the VM-Series firewall are very similar to the hardware versions of PAN-OS. When problems occur, you should check interface counters, system log files, and if necessary, use debug to create captures. For more details on PAN-OS troubleshooting, refer to the Packet Based Troubleshooting Tech Note at *https://live.paloaltonetworks.com/community/documentation*.

Similar to a physical environment, it is sometimes useful to have a separate troubleshooting client to capture traffic in the virtualized environment. It can be helpful to build a fresh OS from scratch with common troubleshooting tools installed, such as: tcpdump, nmap, hping, traceroute, iperf, tcpedit, netcat, etc. Going forward, each time it is needed, the troubleshooting client can be quickly deployed to the virtual switch(es) in question and used to isolate networking problems.

For performance related issues on the firewall, first check the **Dashboard** from the firewall web interface. For information in the VM server, in vSphere Client go to **Home > Inventory > VMs and Templates**, select the VM-Series firewall instance and click the **Summary** tab. Under **Resources**, check the statistics for consumed memory, CPU and storage. For resource history, click the **Performance** tab and monitor resource consumption over time.

You can also view alerts or create a tech support or stats dump files from **Device > Support**. For more information, refer to "Viewing Support Information" on page 117.

Chapter 12

# Setting Up Panorama

This chapter describes how to set up the Panorama centralized management system:


- "Overview" in the next section
- "Setting Up Panorama as a Virtual Appliance" on page 372
- "Setting up Panorama on an M-Series Appliance" on page 376
- "Configuring High Availability (HA)" on page 377

***Note:*** *For information on using Panorama, see "Central Device Management Using Panorama" on page 381.*

## Overview

Panorama is the centralized management system for the Palo Alto Networks family of next-generation firewalls, and is available both as a dedicated hardware platform and as a VMware virtual appliance that meets your server consolidation needs.

Panorama provides centralized visibility and management for all the devices on your network. Because Panorama shares the same web-based look and feel as the individual device interface, you can seamlessly transition in to managing the devices centrally and reducing the administrative effort in managing multiple devices.

For information on installing Panorama on VMware ESX(i) 3.5 or later, see "Setting Up Panorama as a Virtual Appliance" on page 372

For information on setting up the hardware based M-100 management appliance, see "Setting up Panorama on an M-Series Appliance" on page 376.

# Setting Up Panorama as a Virtual Appliance

The pre-requisites for setting up a Panorama virtual appliance are:

- System requirements
  - VMware ESX(i) 3.5 or later
  - Quad Core CPU (3GHz) (use 4 GHz if you have 10 or more active firewalls)
  - 2-4 GB RAM (use 4 GB if you have 10 or more active firewalls)
  - 34 GB disk space
  - VMware vSphere Client 4.x or VMware Infrastructure Client 3.5
- Use the assigned serial number to register Panorama on the support site at *https://support.paloaltonetworks.com*. After you register Panorama on the support site, you will have access to the Panorama software downloads page.
- Download the latest Panorama base image zip file to the server on which you will be installing Panorama.The virtual appliance installation procedure uses the Open Virtual Machine Format (OVF) template file, which is included in the base image.

## Installing Panorama

Follow these steps to install Panorama on your ESX(i) server:

1. Unzip the Panorama zip file to find the *panorama-esx.ovf* template file for installation.
2. Open the VMware vSphere Client and connect to your VMware server from the login screen.
3. Choose **File > Deploy OVF Template**.
4. Browse to select the *panorama-esx.ovf* file from the recently unzipped Panorama base image, and click **Next**.
5. Confirm that the product name and description match the downloaded version, and click **Next**.
6. Choose a name for the Panorama image, and click **Next**.
7. Select a datastore location to install the Panorama image, and click **Next**.
8. If prompted, choose **Thick provisioned format** for the disk format, and click **Next**.
9. Confirm the options you selected and then click **Finish** to begin the installation process.
10. When the installation is complete, choose the newly installed Panorama image and click the **Power On** button.

When the Panorama virtual machine boots, the installation process is complete. Continue with the next section to use the console and perform initial set up.

# Configuring the Panorama Network Interface

To configure the Panorama network interface using the Panorama virtual machine console on your ESX(i) server:

1. Log in to the CLI. Enter **admin** for both the username and password fields.
2. At the prompt, type **configure** to switch to the configuration mode.
3. Define the network access configuration for the management interface. Enter the following commands on one line.
   **set deviceconfig system ip-address** <Panorama-IP> **netmask** <netmask> **default-gateway** <gateway-IP> **dns-setting servers primary** <DNS-IP>

   where <*Panorama-IP*> is the IP address, <*netmask*> is the subnet mask, <*gateway-IP*> is the IP address of the network gateway, and <*DNS-IP*> is the IP address of the Domain Name System (DNS) server.
4. Type **commit** to save the changes and activate them.
5. Type **Exit** to leave the configuration mode.
6. Test network connectivity to your default gateway or another server (<*target-IP*>).
   **ping host <target-IP>**

   Make sure that you can successfully ping the gateway and the Internet.

### Next Steps

- To log in to Panorama, and to change the default password, see "Logging in to Panorama" on page 377.
- To set up additional log storage capacity for your Panorama virtual appliance, see "Expanding the Log Storage Capacity" on page 373.
- To configure high availability, see "Configuring High Availability (HA)" on page 377.
- To start managing devices using Panorama, see the following sections:
  - To add devices, see "Panorama Administrator Roles, Profiles, and Accounts" on page 387.
  - Top verify that each managed device is configured with the IP address of the Panorama server, see "System Setup, Configuration, and License Management" on page 30.

# Expanding the Log Storage Capacity

By default, Panorama maintains internal storage for log files and statistical data. The default Panorama installation is set up with a single disk partition for all data; 10 GB of space is allocated for log storage on the partition. To support environments that require more storage space, you can select from two options:

- Create a custom virtual disk up to 2TB for ESX or ESXi. For instructions, see "Adding a Virtual Disk" on page 374.

  or,

- Configure an external NFS data store. For instructions, see"Setting Up Storage Partitions" on page 374.

# Adding a Virtual Disk

If you require more storage space than the 10GB that is provided by default on the Panorama device, use the following steps to create a custom virtual disk:

1. On your ESX(i) server, select the Panorama virtual machine.
2. Click **Edit Settings**.
3. Click **Add** to launch the Add Hardware wizard.
4. Choose **Hard Disk** from the list of hardware types and click **Next**.
5. Choose the **Create a new virtual disk** option and click **Next**.
6. Choose **SCSI** for the Virtual Disk Type and click **Next**.
7. Select **Specify a datastore** in the location field and enter a name and path or select using the **Browse** button.
8. Click **Finish**.

   The new disk is shown in the list of devices for the virtual machine.
9. Start the Panorama virtual machine.

   On the first start after adding the new disk, Panorama initializes the new disk for use. This process may take several minutes to a few hours, depending on the size of the newly added disk.

After the system starts with the new disk, any existing logs on the default disk are moved to the new virtual disk, and all future log entries are written to the new disk. If the virtual disk is removed, Panorama automatically reverts back to logging to the default internal 10 GB disk.

If you have already added a virtual disk and would like to replace it with a larger or different virtual disk, you must first remove the installed virtual disk. However, when a first virtual disk is removed you can no longer access the logs on that disk.

***Note:*** *To allow for redundancy, use the virtual disk in a RAID configuration. RAID 10 provides the best write performance for applications with high logging characteristics. For further performance improvements, optimize the drives for sequential writing of a small number of large files.*

# Setting Up Storage Partitions

▶ ***Panorama > Setup > Storage Partition Setup***

To configure an external NFS data store:

Click the **Storage Partition Setup** link on the Panorama **Setup** page, and specify the following settings.

***Note:*** *You must reboot the Panorama server after configuring the storage partition settings.*

**Table 144. Storage Partition Settings**

| Field | Description |
|---|---|
| Internal | Maintains storage space for log files and statistical date on the Panorama device. |
| NFS v3 | Specifies an external NFS server mount point for storage. Configure the following settings:<br>• **Server**—Specify the fully qualified domain name (FQDN) or IP address of the NFS server.<br>• **Log Directory**—Specify the full path name of the directory where the logs will be stored.<br>• **Protocol**—Specify the protocol for communication with the NFS server (UDP or TCP).<br>• **Port**—Specify the port for communication with the NFS server.<br>• **Read Size**—Specify the maximum size (bytes) for NFS read operations (range 256 - 32768).<br>• **Write Size**—Specify the maximum size (bytes) for NFS write operations (range 256 - 32768).<br>• **Copy On Setup**—Select the check box to mount the NFS partition and copy any existing logs to the destination directory on the server when the Panorama device boots.<br>• **Test Logging Partition**—Click to perform a test that mounts the NFS partition and presents a success or failure message. |

# Setting up Panorama on an M-Series Appliance

The hardware-based Panorama solution uses the M-100 device. Refer to the *M-100 Hardware Reference Guide* for information on rack-mounting the appliance and powering it on.

Complete the following tasks before you begin with performing initial setup.

- Register your M-100 device at http://support.paloaltonetworks.com to obtain the latest software updates and to activate support.
- Obtain an IP address for Panorama from your system administrator.
- Set your computer's IP address to 192.168.1.2 and the subnet mask to 255.255.255.0.

## Performing Initial Setup

The initial setup allows you to use the default IP address to access the Panorama Management Console and then assign an IP address on the management interface so that you can manage the device across your network.

To perform initial setup:

1. Connect your computer to the management port (MGT) using an RJ-45 Ethernet cable (provided).
2. Power on your computer.
3. Log in to Panorama.
   a. Launch an Internet browser on your computer and enter **https://192.168.1.1**.
   b. Type admin in both the **Name** and **Password** fields.
   c. Click **Login**.
4. Click **Panorama > Setup**. Click the small gear icon (Edit…) on the **Management Interface Settings** table.
5. Enter the new IP address and related network access information for using the management interface (MGT) on your enterprise management network. Click **OK**.
6. Commit your changes to the device. Click **Commit** and select Panorama as the **Commit Type**, then click **OK**.
7. Disconnect your computer from the M-100 device.
8. Connect the MGT port on the front panel of the M-100 to the enterprise management network. Your device is now set up for access across your network.

Next steps:

- To log in and verify access to the Panorama Management Console. See “Logging in to Panorama” on page 377.
- To change the default login credentials, see “Changing the Default Password” on page 377.

# Logging in to Panorama

To log in to the web interface of Panorama:

1. Launch a web browser and enter **https://<Panorama IP address>**.
   The browser automatically opens the Palo Alto Networks login page.
2. Enter **admin** in both the **Name** and **Password** fields.
3. Click **Login**.

## Changing the Default Password

Panorama is configured with a default password. To improve security, you are encouraged to strengthen the password on first-time login.

1. Choose **Panorama > Administrators > admin**.
2. Enter **admin** in the **Old Password** field.
3. Enter a new password (case-sensitive, up to 15 characters) in the **New Password** field and re-enter the password in the **Confirm New Password** field.
4. Click **OK**.
5. Click **Commit** and select **Panorama** as the **Type.**
6. Click **OK**.

# Configuring High Availability (HA)

▶ ***Panorama > High Availability***

***Note:*** *HA is supported only for managed devices running Release 4.0 or later. It is not backward compatible with Release 3.1 or earlier.*

High availability (HA) allows for redundancy in the event of a device failure. For ensuring HA, you can deploy a pair of hardware-based Panorama appliances or a pair of Panorama virtual appliances in a HA peer configuration that provide synchronized connections to the managed firewalls. Among the peers on the HA configuration, one device must be designated as active and the other as passive. The peers maintain a heartbeat, or a periodic ICMP ping, to verify operational status. If the active Panorama device becomes unavailable, the passive server takes over temporarily. With preemption enabled, the default setting, when the active device becomes available again, the passive device relinquishes control and returns to the passive state.

***Note:*** *To configure a HA pair of Panorama virtual appliances, you must have two Panorama licenses with unique serial numbers for each virtual instance.*

When configuring HA for Panorama virtual appliances, you also need to designate a priority level of either primary or secondary for each peer in the pair. This primary or secondary configuration determines which peer is designated as the primary recipient for logs sent by the managed firewalls. You can configure Panorama to use the same log external storage facility for the assigned primary and secondary devices (Network File System or NFS option) or configure logging internally. If you use the NFS option, only the primary recipient receives the logs that are sent from the managed firewalls. However, if local logging is enabled, by default the logs are sent to both the primary and the secondary recipient.

To enable HA on Panorama, configure the followings settings.

**Table 145. Panorama HA Settings**

| Field | Description |
|---|---|
| **Setup** | |
| Enable HA | Select the check box to enable HA. |
| Peer HA IP Address | Enter the IP address of the MGT interface of the Panorama peer. |
| Enable Encryption | Enable encryption after exporting the HA key from the HA peer and importing it onto this device. The HA key on this device must also be exported from this device and imported on the HA peer. When enabled, the MGT interface encrypts communication between the HA peers.<br>The key import/export is done on the **Panorama > Certificate Management > Certificates** page. See "Importing, Exporting and Generating Security Certificates" on page 86.<br>***Note:*** *HA connectivity uses TCP port 28 with encryption enabled and 28769 when encryption is not enabled.* |
| Monitor Hold Time (ms) | Enter the length of time (ms) that the system will wait before acting on a control link failure (1000-60000 ms, default 3000 ms). |
| **Election Settings** | |
| Priority | Choose **Primary** or **Secondary**. |
| Preemptive | Select the check box to enable the primary Panorama device to resume active operation after recovering from a failure. If this setting is off, then the secondary device remains active even after the higher priority device recovers from a failure. |
| Preemption Hold Time (min) | Enter the time a passive device will wait before taking over as the active device (range 1-60 min, default 1). |
| Promotion Hold Time (ms) | Enter the time that the secondary device will wait before taking over (range 0-60000 ms, default 2000). |
| Hello Interval (ms) | Enter the number of milliseconds between the hello packets sent to verify that the other device is operational (ranges 8000-60000 ms, default 8000). |
| Heartbeat Interval (ms) | Specify how frequently Panorama sends ICMP pings to the HA peer (range 1000-60000 ms, default 1000). |
| Monitor Fail Hold Up Time (ms) | Specify the interval that Panorama waits following a path monitor failure before attempting to re-enter the passive state (default 0 ms). During this period, the device is not available to take over for the active device in the event of failure. |
| Additional Master Hold Up Time (ms) | Specify the interval during which the preempting device remains in the passive state before taking over as the active device (default 7000 ms). |

**Table 145. Panorama HA Settings (Continued)**

| Field | Description |
|---|---|
| **Path Monitoring** | |
| Enabled | Select the check box to enable path monitoring. Path monitoring enables Panorama to monitor specified destination IP addresses by sending ICMP ping messages to make sure that they are responsive. |
| Failure Condition | Select whether a failover occurs when any or all of the monitored path groups fail to respond. |
| Path Groups | Define one or more path groups to monitor specific destination addresses. To add a path group, specify the following and click **Add**:<br>• **Name**—Specify a name for the path group.<br>• **Enabled**—Select the check box to enable the path group.<br>• **Failure Condition—**Select whether a failure occurs when any or all of the specified destination addresses fails to respond.<br>• **Ping interval**—Specify a length of time between ICMP echo messages to verify that the path is up (range 1000-60000 ms, default 5000).<br>• **Destination IPs**—Enter one or more destination addresses to be monitored (multiple addresses must be separated by commas).<br>• **Ping Interval**—Specify the interval between pings that are sent to the destination address (range 1000-60000 milliseconds, default 5000 milliseconds).<br>• **Ping Count**—Specify the number of failed pings before declaring a failure (range 3-10 pings, default 3 pings).<br>To delete a path group, select the group, and click **Delete**. |

# Switching the Logging Priority in an HA Pair

When a failover occurs for Panorama virtual appliances in an HA configuration with NFS-based logging, the logging functionality is interrupted. To re-enable logging following a failure, you must promote the secondary Panorama device to function as the primary so that it can establish a connection with the NFS-based log partition.

***Note:*** *For configurations that use internal logging instead of NFS, follow the instructions through Step 2 of the procedure in this section to switch the logging priority of the secondary device to primary.*

For the following procedure, assume that the active primary is running on server S1 and the passive secondary is running on S2. Failover has occurred, and S2 has become the active secondary.

To assign S2 as the primary log recipient, follow these steps:

1. Power S1 off.
2. Configure S2 to be primary and commit the configuration:
   a. Choose **Panorama > High Availability.**
   b. Edit the election settings and change **Priority** from **Secondary** to **Primary**.
   c. Commit the changes, rebooting the device when prompted. The reboot is required because the configuration refers to NFS storage.

3. Execute the CLI command **`request high-availability convert-to-primary`**.

   If S1 is connected as the HA peer to S2 and NFS storage is specified, then the **convert-to-primary** command will fail, indicating that the HA peer (S1) needs to be powered down before the operation can succeed. If the peer is not connected, the system dynamically mounts the NFS disk, converts the ownership of the partition to S2, and unmounts the partition.

4. Reboot S2.

   When S2 comes up, it is able to write to the NFS-based log partition.

# Chapter 13
# Central Device Management Using Panorama

This chapter describes how to use the Panorama centralized management system to manage multiple firewalls:


- "Accessing the Panorama Web Interface" in the next section
- "Using the Panorama Interface" on page 382
- "Adding Devices" on page 385
- "Panorama Administrator Roles, Profiles, and Accounts" on page 387
- "Specifying Panorama Access Domains for Administrators" on page 392
- "Working with Objects" on page 395
- "Working with Policies" on page 393
- "Templates" on page 399
- "Logging" on page 403
- "Viewing Firewall Deployment Information" on page 416
- "Backing Up Firewall Configurations" on page 417
- "Scheduling Configuration Exports" on page 417
- "Upgrading the Panorama Software" on page 419

# Accessing the Panorama Web Interface

To access the Panorama interface for centralized firewall management, log in to the Panorama server web interface:

1. Launch your preferred web browser and enter **https://*Panorama IP address***

   The browser automatically opens the Palo Alto Networks login page.

2. Enter the login name and password and click **Login**.

# Using the Panorama Interface

Panorama allows you to view information about multiple devices in your network and to manage devices from a central web interface.

To display information regarding the Palo Alto Networks firewalls in the network, the devices must be connected to the Panorama server.

Perform these steps to allow the devices to connect:

1. Add the IP address of the Panorama server to each device. Refer to "Defining Management Settings" on page 30.

2. Use the Panorama interface to add the devices. Refer to "Panorama Administrator Roles, Profiles, and Accounts" on page 387.

Certain Panorama configuration tabs will not appear until their respective components are configured. For instance, the Policies and Objects tab will only appear after adding device groups **Panorama > Device Groups** and **Device** and **Network** will only appear after adding templates from P**anorama > Templates**.

The Panorama tabs are described in the following table.

**Table 146. Summary of Panorama Tabs**

| Page | Description |
|---|---|
| Dashboard | Displays general information about the managed devices, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available charts are displayed by default, but each user can remove and add individual charts, as needed. |
| ACC | Displays the overall risk and threat levels for the managed devices. Refer to "Using the Application Command Center" on page 253 and "Identifying Unknown Applications and Taking Action" on page 276. |
| Monitor | Allows you to view logs and reports. Refer to "Viewing Reports" on page 273. |
| Policies | Allows you to define policies to share across managed firewalls. Refer to "Logging and Reporting" on page 403 for information using the pages in this tab. |
| Objects | Allows you to define policy objects that are shared across the managed firewalls. Refer to "Logging and Reporting" on page 403. |
| Network | Allows you to apply network configuration options across managed firewalls using templates. Refer to "Templates" on page 399. |

**Table 146. Summary of Panorama Tabs (Continued)**

| Page | Description |
|---|---|
| Device | Allows you to apply device configuration options across managed firewalls using templates. Refer to "Templates" on page 399. |
| Panorama | Allows you to configure Panorama and manage deployed firewalls. Refer to "Panorama Tab" in the next section. |

# Panorama Tab

The **Panorama** tab is similar to the **Devices** tab for the firewall, but the settings apply to the Panorama device, not the managed firewalls. The following table describes the pages on this tab. To access a page, click the page name link on the side menu.

**Table 147. Summary of Panorama Pages**

| Page | Description |
|---|---|
| Setup | Allows you to specify the Panorama host name, the network settings of the management interface, and the addresses of network servers (DNS and NTP). Refer to "Defining Management Settings" on page 30. |
| Templates | Allows you to create Templates that can be used to manage configuration options based on the Device and Network tabs, enabling you to deploy templates to multiple devices that have similar configurations. Refer to "Templates" on page 399. |
| Config Audit | Allows you to view and compare configuration files. Refer to "Defining Operations Settings" on page 37. |
| Managed Devices | Allows you to add devices for management by Panorama, push shared configuration to managed devices, and run comprehensive configuration audits on devices or entire device groups. Refer to "Adding Devices" on page 385. |
| Device Groups | Allows you to define sets of devices that are treated as a unit when creating objects and applying policies in Panorama. Refer to "Defining Device Groups" on page 386. |
| Managed Collectors | Allows you to configure and manage the log collector devices that will be used in your environment to distribute logging information for firewalls that are managed by Panorama. You can also use this pages to upgrade the software on your log collectors. You first download the latest Panorama software and you can then push the updated version to your log collectors by clicking **Install** on the Managed Collectors page.<br><br>***Note:*** *Log collectors are comprised of the log collector software (part of the Panorama software package) and the M-100 hardware platform. The M-100 can be configured as a Panorama manager, a log collector, or both. The operational command to change the mode of an M-100 is* **`request system logger-mode [panorama | logger]`**.*To view the current mode, run* **`show system info | match logger_mode`**.<br>*When an M-100 is in log collector mode, only the CLI is available for management.*<br><br>Refer to "Managing Log Collectors" on page 407. |

**Table 147. Summary of Panorama Pages (Continued)**

| Page | Description |
|---|---|
| Collector Groups | Allows you to group log collectors together so you can apply the same configuration settings to all collectors in the group. You also use the collector group to assign firewalls to log collectors.<br>***Note:*** *You can add up to 4 log collector devices per collector group. Refer to "Defining Log Collector Groups" on page 411.* |
| Admin Roles | Allows you to specify the privileges and responsibilities that are assigned to users who require access to Panorama. Refer to "Defining Administrator Roles" on page 58. |
| Password Profiles | Allows you to define password profiles, which can then be applied to Panorama administrators. You can configure the following profile options:<br>• Required password change period (days)<br>• Expiration warning period (days)<br>• Post Expiration Admin Login Count<br>• Post Expiration Grace Period (days) |
| Administrators | Allows you to define the accounts for users who require access to Panorama. Refer to "Creating Administrative Accounts" on page 59.<br>***Note:*** *On the* ***Administrators*** *page for "," a lock icon is shown in the right column if an account is locked out. The administrator can click the icon to unlock the account.* |
| High Availability | Allows you to configure a pair of Panorama devices to support high availability (HA). Refer to "Configuring High Availability (HA)" on page 377. |
| Certificate Management | Allows you to configure and manage certificate profiles, trusted certificate authorities, and OCSP responders. Refer to "Importing, Exporting and Generating Security Certificates" on page 86. |
| Log Settings | Allows you to define Simple Network Management Protocol (SNMP) trap sinks, syslog servers, and email addresses for distributing log messages. Refer to "Logging Configuration" on page 70. |
| Server Profiles | Allows you to specify profiles for servers that provide services to Panorama.<br>Refer to the following sections:<br>• "Configuring Email Notification Settings" on page 83<br>• "Configuring SNMP Trap Destinations" on page 75<br>• "Configuring Syslog Servers" on page 76<br>• "Configuring RADIUS Server Settings" on page 65<br>• "Configuring LDAP Server Settings" on page 66<br>• "Configuring Kerberos Settings (Native Active Directory Authentication)" on page 67.<br>• "Configuring Netflow Settings" on page 85 |
| Authentication Profile | Allows you to specify a profile to authentication access to Panorama. Refer to "Authentication Profiles" on page 62. |
| Authentication Sequence | Allows you to specify sets of authentication profiles to use for access to Panorama. Refer to "Authentication Sequence" on page 67. |
| Access Domain | Provides the capability to limit administrators access to Device Groups, Templates, and the device contexts that administrators can switch to based on profiles. Refer to "Panorama Administrator Roles, Profiles, and Accounts" on page 387. |

**Table 147. Summary of Panorama Pages (Continued)**

| Page | Description |
|---|---|
| Scheduled Config Export | Allows you to collect running configurations from managed devices and deliver them daily to a File Transfer Protocol (FTP) server or by using Secure Copy (SCP) to securely transfer data between the Panorama server and a remote host. Refer to "Scheduling Configuration Exports" on page 417. |
| Software | Allows you to view the available Panorama software releases and download and install a selected software version. Refer to "Upgrading the Panorama Software" on page 419. |
| Dynamic Updates | Allows you to view the latest application definitions and information on new security threats, such as antivirus signatures (threat prevention license required) and update Panorama with the new definitions. Refer to "Updating Threat and Application Definitions" on page 55. |
| Support | Allows you to access product and security alerts from Palo Alto Networks. Refer to "Viewing Support Information" on page 117. |
| Deployment | Allows you to view current license information on the managed devices and install software, clients, and dynamic content on the devices. Refer to "Viewing Firewall Deployment Information" on page 416. |
| Master Key and Diagnostics | Allows you to specify a master key to encrypt private keys on the firewall. Private keys are stored in encrypted form by default even if a new master key is not specified. Refer to "Encrypting Private Keys and Passwords on the Firewall" on page 90. |

# Adding Devices

▶ *Panorama > Managed Devices*

The **Managed Devices** page allows you to create a list of devices for centralized management.

If devices are part of an HA pair, you must add both devices or virtual systems of the peers (if in multi-virtual system mode) to the same device group, and Panorama must push the configuration to both HA peer devices at the same time. If you target a rule to specific firewalls that are in an HA configuration, make sure to include both firewalls in the target selection.

***Note:*** *Panorama can manage PAN-OS devices running the same major release or earlier supported versions, but not devices running a later release version. For example, Panorama 4.0 can manage PAN-OS devices running 4.0 or earlier supported versions, but it cannot manage PAN-OS devices running 4.1.*

***Note:*** *Managed devices communicate with Panorama using SSL through TCP port 3978.*

To add devices:

1. Under the **Panorama** tab, click **Managed Devices** to open the **Managed Devices** page.
2. Click **Add** to open an editing window.
3. Enter the serial number of the device to be added, and click **Add**.

4. Add additional devices, as needed.
5. Click **OK**. The window closes and the **Managed Devices** page refreshes to show the newly added devices.
6. To delete a device:
   a. Select the check box next to the device in the Managed Devices table.
   b. Click **Delete**.
   c. Click **OK**.
7. Select the **Group HA Peers** check box to group devices in high availability (HA) mode together.

   This option allows you to easily identify devices that are in HA mode. When pushing shared policies, you can push to the grouped pair, instead of each device individually. Also, when adding a new device in Managed Devices, if they are in HA mode, both devices will be displayed together, so you can add both devices.

   When viewing an HA pair, if the configuration does not match, a warning indicator will appear. You will also see an indicator if the HA devices are in different device groups.

   This option is also independent for each section, so enabling and disabling in one area, will not enable/disable for all areas. The Group HA Peers option is present in the following Panorama areas:

   – Managed Devices
   – Templates
   – Device Groups
   – Policies tab (Target tab for all policy types)
   – Commit dialog

# Defining Device Groups

▶ ***Panorama > Device Groups***

Device groups are used to manage shared policies and objects. You can define device groups that consist of firewalls and/or virtual systems that you want to manage as a group, such as the firewalls that manage a group of branch offices or individual departments in a company. Each group is treated as a single unit when applying policies in Panorama.

You can add each device to at most one device group. Because virtual systems are considered distinct entities in Panorama, you can assign virtual systems within a device to different device groups.

The **Device Groups** page lists the device groups along with the information listed in the following table.

**Table 148. Device Group Settings**

| Field | Description |
|---|---|
| Device Group Name | Enter a name to identify the group (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter a description for the group. |
| Devices | Select the check box next to the device in the Managed Devices table and click **Move,** choose the new group and click **OK**. |
| Master Device | Select a device to use as the master. The master device is the firewall from which Panorama gathers User-ID information for use in policies. The gathered user and group mapping information is specific to a device group and can come from only one device (the master) inside the group. |
| Group HA Peers | Select the check box to group devices in high availability (HA) mode together.<br>This option allows you to easily identify devices that are in HA mode. When pushing shared policies, you can push to the grouped pair, instead of each device individually. Also, when adding a new device in Managed Devices, if they are in HA mode, both devices will be displayed together, so you can add both devices.<br>When viewing an HA pair, if the configuration does not match a warning indicator will appear. You will also see an indicator if the HA devices are in different device groups.<br>This option is also independent for each section, so enabling and disabling in one area, will not enable/disable for all areas. The Group HA Peers option is present in the following Panorama areas:<br>• Managed Devices<br>• Templates<br>• Device Groups<br>• Policies tab (Target tab for all policy types)<br>• Commit dialog |

# Panorama Administrator Roles, Profiles, and Accounts

Panorama supports the following options to authenticate administrative users who attempt to log in to the device:

- **Local database**—The user login and password information is entered directly into the Panorama database.
- **RADIUS**—Existing Remote Authentication Dial In User Service (RADIUS) servers are used to authenticate users.

- **LDAP**—Existing Lightweight Directory Access Protocol (LDAP) servers are used to authenticate users.
- **Kerberos**—Existing Kerberos servers are used to authenticate users.
- **Client Certificate**—Existing client certificates are used to authenticate users.

When you create an administrative account, you specify client certificate (no authentication profile), or an authentication profile (RADIUS, LDAP, Kerberos, or local DB authentication). This setting determines how the administrator password is checked. If you do not specify a profile, the account will use local authentication.

***Note:*** *You may have certain Panorama administrators that may not have access to the* ***Panorama > Administrators*** *menu. In this case, the administrator can click on their username located to the left of the logout link on the bottom of the web interface to change their password.*

Administrator roles determine the functions that the administrator is permitted to perform after logging in. You can assign roles directly to an administrator account, or define role profiles, which specify detailed privileges, and assign those to administrator accounts.

Refer to the following sections for additional information:

- "Setting Up Authentication Profiles" on page 62.
- "Defining Panorama Administrator Roles" on page 388.
- "Certificate Profile" on page 89.
- "Specifying Panorama Access Domains for Administrators" on page 392.
- "Certificate Profile" on page 89.

# Defining Panorama Administrator Roles

▶ ***Panorama > Admin Roles***

Use the **Admin Roles** page to define role profiles that determine the access and responsibilities available to administrative users. For instructions on adding administrator accounts, refer to "Creating Panorama Administrative Accounts" on page 389.

***Note:*** *The Admin Role can be mapped via RADIUS Vendor-Specific Attributes (VSA) using the following attribute: "PaloAlto-Panorama-Admin-Role = <AdminRoleName>,".*

**Table 149. Panorama Administrator Role Settings**

| Field | Description |
|---|---|
| Name | Enter a name to identify this administrator role (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Description | Enter an optional description of the role. |
| Permission | Select the scope of administrative responsibility (Panorama or Device Group and Template). |

**Table 149. Panorama Administrator Role Settings (Continued)**

| Field | Description |
|---|---|
| WebUI | Click the icons for specified areas to indicate the type of access permitted for the web interface:<br>• Read/write access to the indicated page.<br>• Read only access to the indicated page.<br>• No access to the indicated page. |
| XML API | Select the type of access for the XML API<br>• **Report**—Access to the device reports.<br>• **Log**—Access to the device logs.<br>• **Configuration**—Permissions to retrieve or modify the device configuration.<br>• **Operational Requests**—Permissions to run operational commands.<br>• **Commit**—Permissions to commit the configuration.<br>• **User-ID Agent**—Access to the User-ID Agent.<br>• **Export**—Permissions to export files from the device, including the configuration, block or response pages, certificates, keys, and more.<br>• **Import**—Permissions to import files to the device, including software, content, license, configuration, certificates, block pages, custom logs, and more. |
| Command Line | Select the type of role for CLI access:<br>• **None**—Access to the device CLI not permitted.<br>• **superuser**—Full access to the current device.<br>• **superreader**—Read-only access to the current device.<br>• **panorama-admin**—Full access to a selected device, except for defining new accounts or virtual systems. |

# Creating Panorama Administrative Accounts

▶ *Panorama* > *Administrators*

Administrator accounts control access to Panorama. Each administrator can have full or read-only access to Panorama and all managed devices, or can be assigned Panorama administrator access, but will not have permissions to create admin accounts or edit admin roles, which allows access to Panorama configuration and not the managed devices. The predefined **admin** account has full access to Panorama and the managed devices.

The following authentication options are supported:

- Password authentication—The user enters a user name and password to log in. No certificates are required.
- Client certificate authentication (web)—If you select this check box, a user name and password are not required; the certificate is sufficient to authenticate access to the firewall.
- Public key authentication (SSH)—The user can generate a public/private key pair on the machine that requires access to the firewall, and then upload the public key to the firewall to allow secure access without requiring the user to enter a user name and password.

***Note:*** *To ensure that the device management interface remains secure, it is recommended that administrative change their passwords periodically using a mixture of lower-case letters, upper-case letters, and numbers. You can also enforce "Minimum Password Complexity" from Setup > Management.*

**Table 150. Administrator Account Settings**

| Field | Description |
|---|---|
| Name | Enter a login name for the user (up to 15 characters). The name is case-sensitive and must be unique. Use only letters, numbers, hyphens, and underscores. |
| Authentication Profile | Select an authentication profile for administrator authentication according to the settings in the specified authentication profile. This setting can be used for RADIUS, LDAP, or Kerberos authentication.<br>For instructions on setting up authentication profiles, refer to "Setting Up Authentication Profiles" on page 62. |
| Use only client certificate authentication (Web) | Select the check box to use client certificate authentication for web access. If you select this check box, a user name and password are not required; the certificate is sufficient to authenticate access Panorama. |
| Password/Confirm Password | Enter and confirm a case-sensitive password for the user (up to 15 characters). *You can also enforce "Minimum Password Complexity" from Setup > Management.*<br>You may have certain Panorama administrators that do not have access to the Panorama > Administrators menu, in this case, the administrator can click on their username located to the left of the logout link on the bottom of the web interface and change their local password from there. |

**Table 150. Administrator Account Settings (Continued)**

| Field | Description |
|---|---|
| Use Public Key Authentication (SSH) | Select the check box to use SSH public key authentication. Click **Import Key** and browse to select the public key file. The uploaded key is displayed in the read-only text area.<br>Supported key file formats are IETF SECSH and OpenSSH. Supported key algorithms are DSA (1024 bits) and RSA (768-4096 bits).<br>***Note:*** *If the public key authentication fails, a login and password prompt is presented to the user.* |
| Role | Select an option for assigning a role to this user. The role determines what the user can view and modify.<br>• **Dynamic**, you can select any of the following pre-specified roles from the drop-down list:<br>– **Superuser—**Full access to the current device.<br>– **Superuser (Read Only)—**Read-only access to the current device.<br>– **Panorama administrator—**Full access to the Panorama instances.<br>• **Role Based —**Access based on assigned roles, as defined in "Defining Panorama Administrator Roles" on page 388.<br>If you choose Role Based, select a previously-defined role profile from the drop-down list. For instructions on defining role profiles, refer to "Defining Panorama Administrator Roles" on page 388<br>For role based access, when you select a profile assigned to the **Device Group and Template** administrator role, the **Access Control** tab is displayed. On this **Access Control** tab, you can define access to Device Groups, Templates, and Device Context. The definitions for these fields are the same as the information in "Specifying Panorama Access Domains for Administrators" on page 392. |

***Note:*** *On the Panorama* ***Administrators*** *page for "," a lock icon is shown in the right column if an account is locked out. The administrator can click the icon to unlock the account.*

# Specifying Panorama Access Domains for Administrators

▶ *Panorama > Access Domain*

Use the **Access Domain** page to specify domains for role-based administrators who have access to device groups and templates. Adding a device group to an access domain allows you to manage policies and objects for that device group. Adding an individual firewall to an access domain allows you to switch into the device context for that firewall.

The access domain is linked to RADIUS vendor-specific attributes (VSAs) and is supported only if a RADIUS server is used for administrator authentication. If RADIUS is not used, the access domain settings on this page are ignored.

***Note:*** *The Access Domain can be mapped via RADIUS VSA using the following attribute: "PaloAlto-Panorama-Admin-Access-Domain = <AccessDomainName>,".*

**Table 151. Access Domain Settings**

| Field | Description |
|---|---|
| Name | Enter a name for the access domain (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, hyphens, and underscores. |
| Device Groups | Click **Add** to specify pre-defined device groups to include in the access domain. |
| Device Context | Select the device(s) that the administrator can do a context switch to in order to allow local configuration edits. |
| Templates | Click **Add** to specify pre-defined templates to include in the access domain. |

# Device Groups

Panorama device groups allow you to group firewall devices and then define policies and objects that can be shared across those device groups.

The following sections describe how to define policies and objects for device groups:

- "Working with Policies" in the next section
- "Working with Objects" on page 395

## Working with Policies

▶ *Policies*

Panorama allows you to define policies that are shared across the managed firewalls. You can apply pre and post rules that will apply to a device group and you can add additional global pre and post rules that apply to all device groups. This creates a layered approach of applying policies to managed devices. The first layer is the device level rules that are local to the device, you can then apply pre and post rules to device groups, and then you can add another layer of global pre and post rules that apply to all device group in the Panorama instance as shown in Figure 47.

**Figure 47. Shared Policy Ordering**

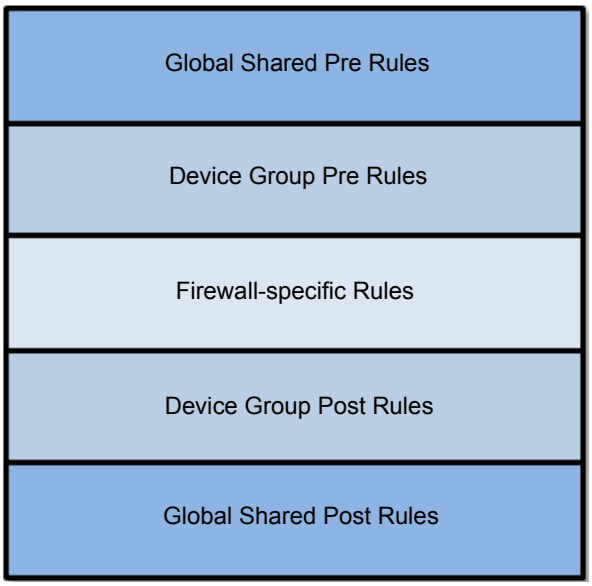


General information about working with policies is found in "Policies" on page 183. This section describes the modifications and best practices that apply to policies in Panorama.

The following best practices apply to policies in Panorama:

- **Pre rules:** Pre rules are evaluated before any device specific rules and generally make up the majority of a deployment's shared rulebase. Do not add any pre rules if you will need device-level exceptions.

  If you do not want administrators to be able to allow any applications at specific sites, you can include a deny rule for all zones, users, and applications as your last rule in the set of pre rules.

- **Firewall-specific rules:** Define rules for an individual firewall to create site-specific policies.
- **Post rules**: Use these rules to specify what happens to traffic that is not covered by the pre rules and firewall-specific rules. For example, if a pre rule specifies certain allowed applications and the post rule is "deny all," then applications not covered by the pre rule are stopped. You can then add rules to allow additional applications based on user request. You can also create device level "allow" rules as exceptions for specific applications that are permitted at an individual location.
- **Global pre rules:** Global pre rules are evaluated before any firewall-specific, or device group rules and are applied to all devices in managed device groups.
- **Global post rules:** Use these rules to specify what happens to traffic that is not covered by device group pre rules and firewall-specific rules. These rules are evaluated last, and only after the device group pre rules.

  Global shared rules can only be created and modified by the Panorama Admin or superuser. So these rules can be used to apply policies before and after any rules applied by device group administrators.

The following applies when defining policies in Panorama:

Panorama applies policies to specified device groups and global shared policies applies another layer of policies to the device groups.

When you create a device group rule from the Panorama instance, you select the **Policies** tab, choose the device group that you want to create rules for and then define your rules. To create global rules that apply to all managed devices in device groups, you choose the **Shared** option in the **Device Group** drop down, as shown in Figure 48.

**Figure 48. Panorama Shared Policy**

You can also specify that shared objects take precedence over device group objects by selecting the check box **Shared Objects Take Precedence** in **Panorama > Setup > Management > Panorama Settings**. This option is a system-wide setting and is off by default. When this option is off, device groups override corresponding objects of the same name. If the option is on (checked), device group objects cannot override corresponding objects of the same name from a shared location and any device group object with the same name as a shared object will be discarded.

- You can target a policy rule to individual devices within the device group for which the rule is defined. To target a device after a policy is created, click an entry in the **Target** column and select the devices in the pop-up window. If you do not select devices to target, the policy is added to all of the (unchecked) devices in the device group.

  To apply the rule to all devices in a device group EXCEPT the targeted device, select the **Install on all but specified devices** check box and select the devices that you want to exclude. If you do not select any device, the policy is added to none of the devices in the device group.

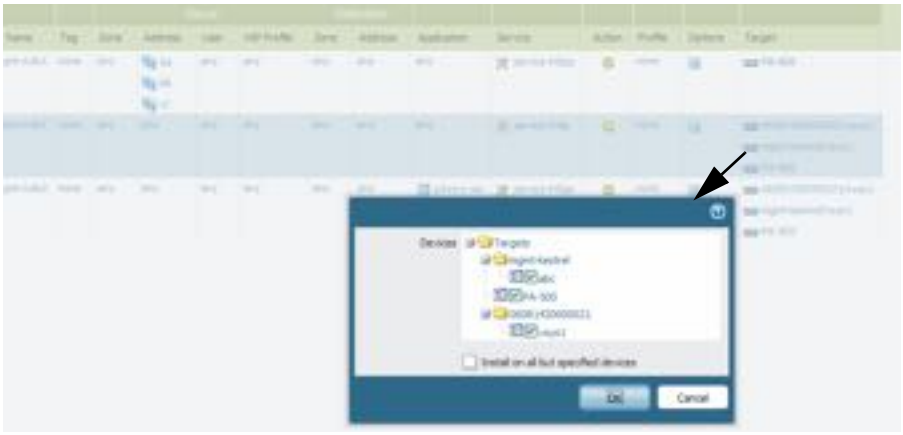

**Figure 49. Targeting Policy Rules to Individual Devices in Panorama**

- Zones are not created in Panorama, but you can select zones based on zone names that are gathered from any template which maps to the same devices as the ones included in a device group. If a zone is not available from the templates, you must manually type a zone name when you first create a rule. For subsequent rules, you can enter new zones or select from previously entered zones.

***Note:*** *The gathering of objects from templates not only applies to zones, it applies to all objects that can be referenced from the Policy/Objects tab, which comes from the Device/Network tabs. This includes: zones, interfaces, certificates, Network Profiles > Monitor, server profiles (Syslog, Email, SNMP Traps), virtual systems in PBF rules, and the local user databases for users/groups.*

- Each policy type listed on the side menu includes pages to define pre rules and post rules as well as global shared rules that are applied on top of the device group pre and post rules. See page 393 for information on best practices when using policies.

# Working with Objects

Panorama supports sharing of objects defined in Panorama. You can create objects on Panorama and then push the object configurations to the managed firewalls. The objects become available for use in policies that are defined on the individual managed firewalls.

***Note:*** *All custom objects should have unique names, and predefined names such as "any" or "default" should be avoided. In particular, using the same object names with different device groups can cause confusion on devices and on Panorama.*

All objects in the **Objects** tab and some objects in the **Device** tab can be managed centrally. **Device** tab objects are managed under the **Panorama** tab and include the following: certificates, response pages, server profiles (SNMP trap, syslog, email, RADIUS, LDAP, and Kerberos), authentication profiles and sequences, and certificate profiles. These objects have a **Location** field that allows you to select where the object should exist in the deployment (for example, "device-group-test"). The following table explains the available object assignment and sharing options for the **Location** field.

**Table 152. Object Assignment and Sharing Options**

| Field | Description |
|---|---|
| Panorama | Panorama supports keeping objects locally and not pushing the objects to any managed devices. To do this, choose **Panorama** from the **Location** drop-down list when defining the object.<br>This option is available only on the **Panorama** tab and only for the following:<br>• Server profiles, including SNMP Trap, Syslog, Email, Remote Authentication Dial In User Service (RADIUS), Lightweight Directory Access Protocol (LDAP), and Kerberos<br>• Authentication profiles, authentication sequences, and certificate profiles |
| Device Groups | Device groups are used to make objects and policies defined in Panorama available to specified sets of devices. For information on creating device groups, see "Defining Device Groups" on page 386.<br>• On the **Policies** or **Objects** tab, choose the device group from the **Device Groups** drop-down list when defining the object.<br>***Note:*** *If you have objects of the same name where one is shared and another is device group specific, the device group specific object will be used for that device group.*<br> 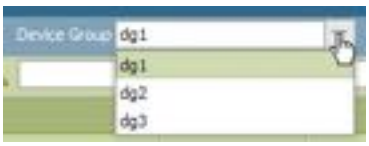 <br>• On the **Panorama** tab, choose a device group from the **Location** drop-down list when defining the object. |
| Shared | Creating a shared object makes the object available for use in any device group. Only Panorama and administrators can create objects in the shared location.<br>• On the **Panorama** tab, choose **Shared** from the **Location** drop-down list when defining the object.<br>• On the **Objects** tab, select the **Shared** check box when defining the object. |

# Working with Devices

Switching context allows an administrator to switch from managing shared policy on Panorama to managing device-specific settings on an individual firewall (such as device specific policy, networking, and device setup). Use the **Context** drop-down list above the side menu to choose an individual device or the full Panorama view. You can select the name of any device that has been added for management by Panorama (refer to "Panorama Administrator Roles, Profiles, and Accounts" on page 387). When you select a device, the web interface refreshes to show all the device tabs and options, allowing you to manage all aspects of the device from Panorama.

***Note:*** *You can only switch context to connected devices. Disconnected devices are not shown in the drop-down list.*

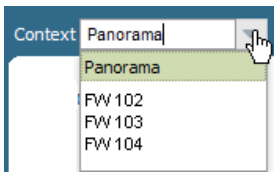


**Figure 50. Choosing Context**

# Commit Operation in Panorama

To commit Panorama configuration changes click the **Commit** icon to bring up the commit dialog box. This dialog box allows you to commit specific areas of the Panorama environment, see Figure 51.

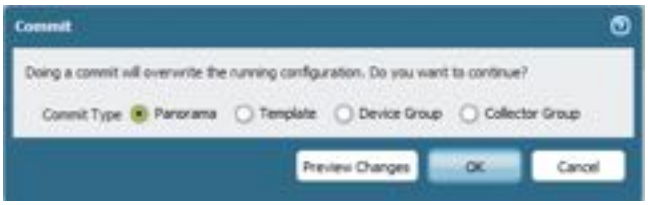


**Figure 51. Panorama Commit Dialog Box**

The following options are available in the commit dialog box:

*Note: When performing commit operations in Panorama, it is suggested that you commit the Panorama configuration first, before committing configuration updates to managed devices.*

- **Commit Type**—Choose the commit type:
  - **Panorama**—Commit the current candidate configuration for Panorama.
  - **Template**—Commit template changes from Panorama to the selected devices. When committing templates, you can select a subset of devices if desired.
  - **Device Group**—Commit device configuration changes from Panorama to the selected device/virtual system(s).
  - **Collector Groups**—Only commit changes to Panorama log collector groups. This will commit changes made in the **Panorama > Collector Groups** page and will apply those changes to the log collector device.
- **Include Device and Network Templates**—This option is available when committing a Device Group from Panorama and is a combo operation that will include both the device and network template changes. The template that will be applied to the device is the template that the device belongs to as defined in Panorama > Templates. You can also select Commit Type Template to commit templates to devices.
- **Force Template Values**—When doing a Commit Type Template, you can select this option to remove objects on the selected devices or virtual systems that have been overridden by the local configuration. When doing a Commit Type "Device Group", you need to also select the "Include Device and Network Templates" check box since overriding can only occur for Template pushed configuration options. This will cause the overridden objects to inherit settings from the template. Refer to "Overriding Template Settings" on page 402.
- **Merge with Candidate Config**—Choose this option to cause the device to include its local candidate configuration when the commit is invoked from Panorama. If this option is not checked, the device local candidate config is not included.

It is important to leave this option unchecked when you have local administrators making changes on a device and you don't want to include their changes when pushing a configuration from Panorama.

– **Preview Changes**—Click this button to bring up a configuration audit window that shows proposed changes in the candidate configuration compared to the current running configuration. You can choose the number of lines of context to display, or show all lines based on items that have been added, modified, or deleted. This option is available when using Commit Type Device Group, Template, or Panorama.
You can also view the commit state for devices by navigating to **Panorama > Managed Devices** and viewing the **Last Commit State** column.
The **Device > Config Audit** feature performs the same function, refer to "Comparing Configuration Files" on page 49.

After the commit is complete, you will see a "Commit succeeded" messages, if there are warning messages, you will see "Commit succeeded with warnings". To view warnings, navigate to **Panorama > Managed Devices** and see the **Last Commit State** column and click the text to view details.

## Panorama Backward Compatibility

When upgrading Panorama to 5.0, and you are managing devices with PAN-OS 4.0 or earlier, simple style vulnerability protection and anti-spyware profiles are automatically converted to rules of equivalent meaning on those devices. Custom style profiles are converted to exceptions that specify signature-specific actions, with no rules required. After migration, a limited set of changes can be made to the migrated profiles in Panorama if compatibility with devices running PAN-OS 4.0 and earlier is required.

For rules created during conversion from a simple style profile, the action of the migrated rules can be modified, and additional allow exceptions can be added to the exceptions list. If a custom style profile was converted to an exceptions-based profile, the exceptions list can be freely modified, but no rules can be created. If the administrator attempts to commit using an incompatible profile, the commit will fail and the failure will be noted in the Managed Devices list under the Last Commit All State column.

To manage devices running different versions of software you must select, in the Device Groups configuration, a master device which is running the lowest software version of the User-ID agent to remain backward compatible and be able to push user based rules to older firewalls.

# Templates

▶ ***Panorama > Templates***

The Panorama **Templates** page is used to create templates that can be used to manage configuration options based on the Device and Network tabs, enabling you to deploy configurations to multiple devices that require similar settings. You can also deploy a base configuration and if needed, override specific settings on the device. For example, you can deploy a base config to a global group of devices, but configure specific time zones settings directly on the devices based on their location.

When managing device configuration with Panorama, you can use a combination of **Device Group** configuration (to manage shared policies and objects) settings and **Templates** settings (to apply device and network settings), but these features are managed separately because of the differences in what can be configured. For information on adding and configuring Panorama templates, refer to "Configuring Panorama Templates" on page 401.

**Table 153 Template Settings (Panorama)**

| Field | Description |
|---|---|
| **Name** | Enter a template name (up to 31 characters). Use only letters, numbers, spaces, hyphens, periods, and underscores. The name is case-sensitive and must be unique.<br>This name will appear in the **Device** and **Network** tab in the **Template** drop-down menu. When selecting a template from one of these tabs, the settings that are modified will only apply to the selected template. |
| **Description** | Enter a description for the template. |
| **Virtual Systems**<br><br>**Note:** You cannot use a template to create virtual systems on the device. | Select the check box if the template will be used on devices with multiple virtual systems. When defining template settings for multi-virtual system devices, you need to configure settings for each virtual system on the device.<br>**Note:** A template enabled for devices with multi-virtual systems (multi-vsys) cannot be pushed to devices with a single virtual system. When you upgrade your Panorama server to v5.0 and later, by default templates are created for configurations relating to the network and device tabs. If, for example, you have defined a server profile for the managed firewall(s), a template is automatically generated for the server profile. This auto-generated template is enabled for multiple virtual systems. To prevent a commit failure, make sure to clear the **Virtual Systems** checkbox before you push the template to a device that is not multi-vsys capable or if the multi- vsys capability is disabled on the device. |
| **Operational Mode** | Set the PAN-OS operational mode for the template: normal, fips, or cc. |
| **VPN Disable Mode** | Selecting this check box will hide all VPN related options in the **Device** and **Network** tabs. The ability to install GlobalProtect Portal or Gateway licenses is also disabled in this mode.<br>**Note:** This option is designed for countries that do not allow VPN connectivity. Palo Alto Networks hardware models that have the -NV indicator in the model name are hard coded to not allow VPN configurations, so this option should be used when creating templates for these models. |

**Table 153 Template Settings (Panorama)**

| Field | Description |
|---|---|
| **Devices** | This window lists all devices managed by Panorama and will also show device groups. Click the check box next to devices to make them a member of the template. You can also add all devices in a device group by clicking the check box next to the device group.<br>**Note:** Templates are based on individual devices, not device groups. Adding a new device to a device group will not automatically add it to a template. |
| **Group HA Peers** | Select the check box to group devices in high availability (HA) mode together.<br>This option allows you to easily identify devices that are in HA mode. When pushing shared policies, you can push to the grouped pair, instead of each device individually. Also, when adding a new device in Managed Devices, if they are in HA mode, both devices will be displayed together, so you can add both devices.<br>When viewing an HA pair, if the configuration does not match a warning indicator will appear. You will also see an indicator if the HA devices are in different device groups.<br>This option is also independent for each section, so enabling and disabling in one area, will not enable/disable for all areas. The Group HA Peers option is present in the following Panorama areas:<br>• Managed Devices<br>• Templates<br>• Device Groups<br>• Policies tab (Target tab for all policy types)<br>• Commit dialog |

# Configuring Panorama Templates

▶ *Panorama > Templates*

To configure Panorama templates, you first create the template and then add devices to it. After the first template is created, you will then see a **Template** drop-down menu in the **Device** and **Network** tabs. Select the desired template from the **Template** drop-down menu and configure device and network settings as if you were managing a single device, but all options that are set will only be applied to the selected template. After the template is configured, you can do a commit from Panorama that only applies to templates.

To create and configure a template, follow these steps:

## Adding a New Template

1. Under the **Panorama** tab, click **Templates** to open the templates page.
2. Click **Add** and enter the template configuration options. Refer to "Template Settings (Panorama)" on page 400.
3. In the **Devices** field, you will see a list of all devices that are managed by Panorama. Click the check box next to each item to make them a member of the new template. If you select a device group, all devices in that group will be selected.
4. Click **OK** to save the template.

## Configuring a Template

1. Now that a template has been created, click the **Device** or **Network** tab and you will see a **Template** drop-down menu, as shown in Figure 52.

**Figure 52. Template Menu**

2. Click the **Template** drop-down and select the template that you want to configure.
3. Click the **Device** or **Network** tab and set the desired configuration options for the template.

***Note:*** *When a template is selected from the Device or Network tabs, you can only set options that are part of the configuration. You can not set operational type options such as changing the device to multi-virtual system mode, and setting a Master Key.*

4. After making all of the configuration changes, click **Commit** and from the **Commit Type** drop-down menu select **Template**. You can also use the **Device Group** commit option and select **Include Device and Network Templates** check box to push templates to a device group.
5. Click the check box next to each template that you want to commit and then click **OK**. You can also preview your changes from the **Commit** window by clicking the **Preview Changes** button. A pop-up will appear showing commit status.

## Overriding Template Settings

When you apply a template to control device and network settings on a firewall, you may want to override some of those settings and have them controlled by the local device configuration. Example, you can deploy a base config to a global group of devices, but configure specific time zones settings directly on the devices based on their location using an override.

To override device and network setting applied by a template, you simply change to the device context, or access the device directly, navigate to the desired setting and then click the **Override** button. The setting will be copied to the local configuration of the device and will no longer be controlled by the template. You can also revert the change by clicking the **Restore** button and the setting will once again be inherited from the template. When doing a commit from Panorama to a managed device that contains overrides, you can select the **Force Template Values** check box to have Panorama templates take over any overridden objects.

When overriding **Device > Setup and Device > High Availability** settings, the overrides are for individual values and parameters inside of configuration trees, and are not applied to an entire tree configuration. This includes items such as DNS servers, Management IP, or NTP server settings. For items such as interfaces and RADIUS server profiles, you apply overrides to the entire object, not internal values.

To identify settings that have templates applied, you will see the following indicators as shown in Figure 53:

**Figure 53. Template Indicators**

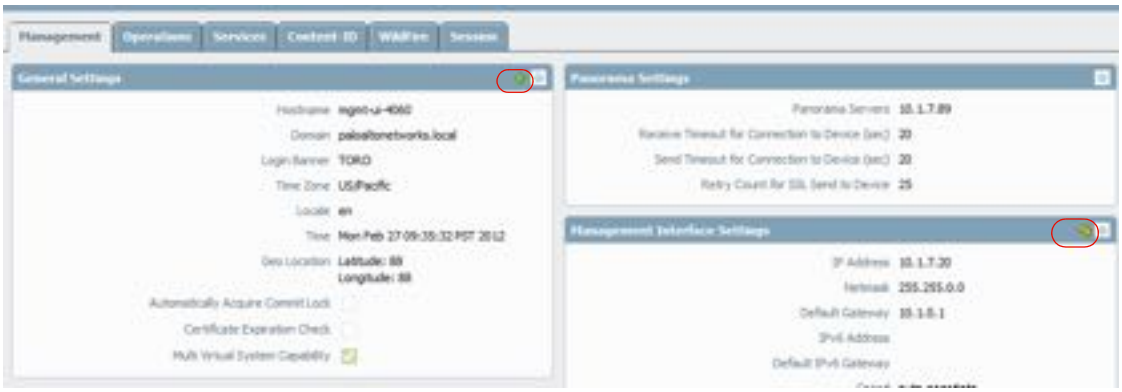

The single green icon indicates that a templates has been applied and there are no overrides. The green and orange icon indicates that a template has been applied and some settings have been overridden.

### Removing Templates

To remove a template, you must disable the template on the local device. On the managed device, navigate to **Device > Setup > Management** tab, then edit the **Panorama Settings** page and click the **Disable Device and Network Template** button. Removing the device from the configuration in **Panorama > Templates** will not delete the template values on the local device.

# Logging

Panorama performs two functions: device management and log collection. To facilitate scalability in large deployments, you can use the M-100 appliance to separate the management and log collection functions on Panorama.

The following sections describe options available for log collection:

- "Logging and Reporting" in the next section
- "Using Panorama for Log Collection" on page 404

## Logging and Reporting

The Panorama logs and reports provide information about user activity in the managed network. Report statistics are aggregated every 15 minutes for use in scheduled predefined and custom reports and statistics are forwarded to Panorama on an hourly basis. If log forwarding is enabled, logs are sent when they are generated on the device.

The **ACC** tab in Panorama displays information from the connected firewalls, all the tables pull information dynamically from the firewalls; they do not require explicit log forwarding. Log forwarding is required for long term log storage and for reporting on logs stored locally in Panorama.

# Generating User Activity Reports

▶ ***Monitor > PDF Reports > User Activity Report***

The Panorama user activity report summarizes user activity across all of the managed firewalls. It is based on firewall data that has been forwarded to Panorama. Refer to "Managing User Activity Reports" on page 272 for general information on creating user activity reports.

# Using Panorama for Log Collection

The M-100 appliance provides a comprehensive log collection solution for Palo Alto Networks firewalls. It helps offload the intensive log collection process from your Panorama management server. Once deployed, each firewall in your managed environment can be configured to send logs to one or more log collectors. By configuring multiple log collectors per firewall, storage capacity is spread out among the collectors creating a more flexible environment with redundancy.

The ACC, PDF Reports, and Logs viewer on Panorama are used to query aggregated information for all managed firewalls. This is the same process that is used if logs were sent directly to Panorama, but in this case, Panorama queries the log collectors to gather the information.

For information on the M-100 hardware, refer to the *M-100 Hardware Reference Guide.*

***Note:*** *The M-100 appliance is shipped with Panorama pre-loaded and both Panorama management and log collection functions are enabled by default. However, if you plan to use the M-100 for log collection only, you need to follow these steps:*

1. *Perform initial configuration to set up an IP address for the management interface and change the default password on the appliance.*
2. *Apply the license on the appliance.*
3. *Verify that the firewalls on your network are running PAN-OS 5.0. If not, you must first upgrade to you PAN-OS 5.0 because support for separate appliances for Panorama management and log collection is not available on earlier versions of PAN-OS*
4. *Change the operational mode on the appliance to log collector mode. For details on the CLI commands, see "Configuring the M-100 as a Log Collector" on page 405.*
   *When the M-100 is in log collector mode, there is no Web interface; Only CLI is available for managing/configuring the appliance.*

The following sections describe how to deploy log collectors:

- "Deploying Distributed Log Collection" in the next section
- "Configuring the M-100 as a Log Collector" on page 405
- "Configuring the Panorama Server for Log Collector Management" on page 406
- "Managing Log Collectors" on page 407
- "Defining Log Collector Groups" on page 411
- "Log Collector Storage" on page 414

## Deploying Distributed Log Collection

Several things should be considered when deploying a distributed log collection solution, including:

- Determine the location of the M-100 appliances based on your network topology and the location of your firewalls. Ideally, the Panorama management server, firewalls, and log collectors should all be connected over a management network. You can configure the firewalls to connect to multiple log collectors in case of failures, so its important that each firewall can reach its assigned collector(s). Typically, the Panorama server and log collectors would be installed in data centers.
- Determine the disk storage needs for log collection based on the number of firewalls that it will handle and the desired retention time. You should analyze your current environment to determine the amount of logs and reports that are being generated.

The illustration in Figure 54 shows a basic log collection deployment. Managed firewalls are configured to send log information to the M-100 log collector. The Panorama servers in HA mode then communicate with the log collectors to perform reporting and to view log information collected from the managed firewalls.

In large deployments where multiple log collectors are available, you assign two or more log collectors to each firewall in a preference order. The first log collector you specify will be the primary log collector for the firewall. If the primary log collector fails, the firewall will momentarily cache log information, so nothing is lost, and will then send logs to the secondary log collector. Regardless of firewall to log collector assignments, logs are balanced across all log collectors in the group to maintain even storage usage among the collector group.

**Figure 54. Distributed Log Collection**

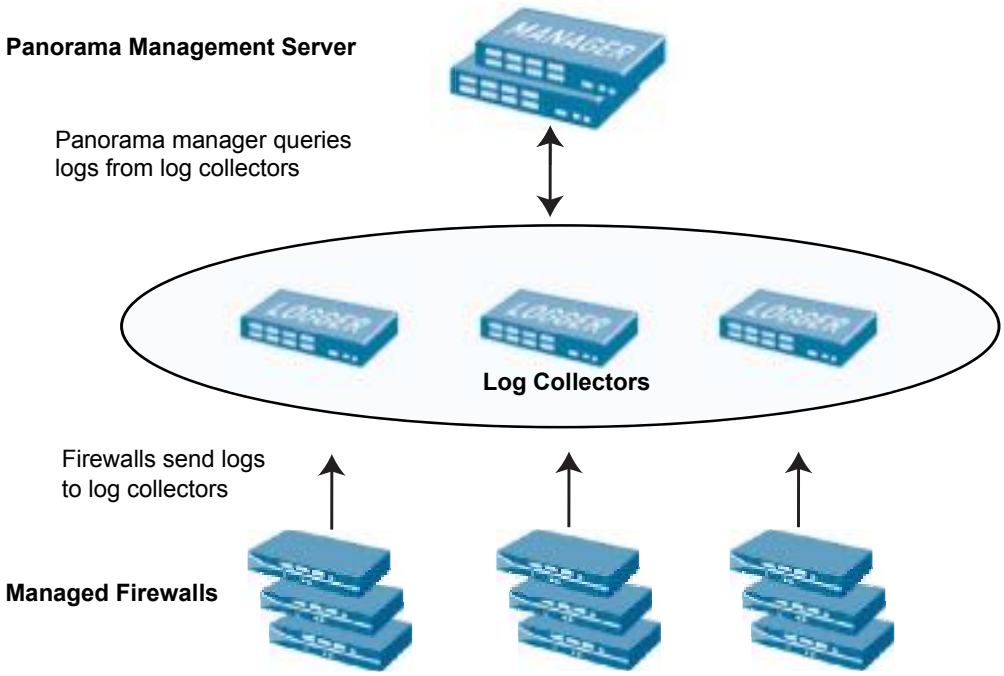


### Configuring the M-100 as a Log Collector

To configure the M-100 appliance as a log collector, follow these steps:

1. Perform initial configuration to set up an IP address for the management interface (MGT port) and change the default password on the appliance. The MGT port is used for all communications, including device management, log collection, and log collector to log collector communications.
2. Apply a license to the M-100.
3. Before you proceed, verify that the firewalls on your network are running PAN-OS 5.0. If not, you must first upgrade the firewalls to PAN-OS 5.0. This upgrade is necessary because support for separate appliances that perform the Panorama management and log collection functionality is not available on earlier versions of PAN-OS.
4. Convert the appliance to function as a log collector only. To change the functionality of the appliance to logger-mode:
   a. Log in to the CLI. For instructions on accessing the CLI, refer to the *Command Line Interface Reference Guide*.
   b. Enter the following CLI command:

      **`request system logger-mode logger`**

      Answer **Yes** to confirm the change. The device will reboot.
   c. After the appliance boots up, verify that the device is in logger mode, using the CLI command:

      **`show system info | match logger_mode`**

      The output should display: **`logger-mode: True`**

***Note:*** *When the M-100 is in log collector mode, only the CLI is available for management. Most of the configuration is performed on the Panorama manager.*

5. Configure the IP address or FQDN of the Panorama server that will be used to manage the log collector.
   `set deviceconfig system panorama-server ip-address <Panorama-IP>.`
   If your Panorama server is in HA mode, enter the IP address of the peer device:
   `set deviceconfig system panorama-server-2 ip-address <Panorama-IP>.`
6. Type **commit** to make the change active and then **exit** to leave configuration mode.

Now that the log collector has the basic configuration needed, for further configuration on the log-collector, use the Panorama Management server to complete the configuration. Refer to "Configuring the Panorama Server for Log Collector Management" in the next section.

### Configuring the Panorama Server for Log Collector Management

The Panorama management server is used to configure and manage the M-100 log collector once network connectivity is established between the two systems. From Panorama, you manage all of the log collector settings, storage, log retention, SNMP settings, and the assignment of managed firewall to log collectors using collector groups. Once the log collectors are in place, Panorama then queries those log collectors for aggregated reporting. Refer to "Defining Log Collector Groups" on page 411.

***Note:*** *You cannot migrate existing firewall logs to the M-100; only new log data will be forwarded after the configuration is completed.*

For detailed descriptions of each of the Managed Collectors fields, refer to "Managing Log Collectors" on page 407.

To configure log collectors from Panorama:

1. Make sure the log collector you want to manage has basic network connectivity so the Panorama management server can communicate with it. Refer to "Configuring the M-100 as a Log Collector" on page 405.
2. From the Panorama management server, navigate to **Panorama > Managed Collectors** and then click **Add** to start adding a log collector.
3. In the **General** tab, specify the log collector serial number and log collector name (hostname). Then specify the Panorama server IP address, DNS, NTP, Timezone, and location information.
4. Use the **Authentication** tab to update the local admin account password on the log collector. You can only use a password hash for this field. You can create a password hash using the Panorama CLI command **`request password-hash password`** ***`password`***. Press enter and the hash value will be displayed. Copy the hash and paste into the **Password Hash** field and **Confirm Password hash** field.
5. The **Management** tab is used to configure the management port on the log collector and is labeled MGT on the front of the M-100. You can set the MTU and other interface settings, the allowed services and define a list of permitted IP addresses that will be allowed to manage the log collector. The MGT port will be used for all communications between the log collector and managed devices. Refer to "Defining Log Collector Groups" on page 411.
6. If you are increasing storage capacity, see "Log Collector Storage" on page 414 for information on configuring the additional disks as a RAID pair.
7. Click **Commit** and in the **Commit Type** drop down select **Panorama** and click **OK**.
8. Click **Commit** and in the **Commit Type** drop down select **Collector Group**, select the log collector or group that you would like to commit and then click **OK**. This will push the changes to the selected log collector(s).
   You will see a pop-up window that will show the commit state.
9. The log collector should now be configured. The next procedure is to create **Collector Groups** to assign firewalls to log collectors. Refer to "Defining Log Collector Groups" on page 411.

## Managing Log Collectors

▶ ***Panorama > Managed Collectors***

Use the Managed Collectors page to configure, manage, and update log collector devices. The settings on this page also exist in the log collector CLI.

After you add log collectors, you can click the **Statistics** link for each collector. This will show the **Collector Statistics** window where you can view Disk Information, Performance numbers for the CPU and Average Logs/sec. You can also view information on the Oldest log received by the collector to get a better understanding of the log range you are looking at.

**Table 154 Managed Collectors Page**

| Field | Description |
|---|---|
| **General Tab** | |
| Collector S/N | Enter the serial number of the log collector device. |
| Collector Name | Enter a name to identify this log collector (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.<br>This is the hostname of the log collector. |
| Panorama Server IP | Specify the IP address of the Panorama server used to manage this collector. |
| Panorama Server IP 2 | Specify the IP address of the secondary device if the Panorama management server is in HA mode. |
| Primary DNS Server | Enter the IP address of the primary DNS server. The server is used for DNS queries from the log collector, for example, to find the Panorama server. |
| Secondary DNS Server | Enter the IP address of a secondary DNS server to use if the primary server is unavailable (optional). |
| Primary NTP Server | Enter the IP address or host name of the primary NTP server, if any. If you do not use NTP servers, you can set the device time manually. |
| Secondary NTP Server | Enter the IP address or host name of secondary NTP servers to use if the primary server is unavailable (optional). |
| Timezone | Select the time zone of the log collector. |
| Latitude | Enter the latitude (-90.0 to 90.0) of the log collector that is used in the traffic and threat maps for App-Scope. |
| Longitude | Enter the longitude (-180.0 to 180.0) of the log collector that is used in the traffic and threat maps for App-Scope. |
| **Authentication Tab** | |
| User Name | This field will always show **admin** and is used for the local CLI login name on the log collector. |

**Table 154 Managed Collectors Page**

| Field | Description |
|---|---|
| Mode | Select **Password** to manually enter and confirm a password to be used for local CLI authentication, or select **Password Hash** to enter a hash value.<br>To create a password hash from the Panorama management server CLI, run the following:<br>`request password-hash password password123`<br>This will return a hash value for the password *password123 (Example- $1$urlishri$aLP2by.u2A1IQ/Njh5TFy9).*<br>Copy the hash value from the CLI and paste to the **Password Hash** field. When you commit your changes, the new hash will be pushed to the log collector and the new local admin login password will be *password123*. |
| Failed Attempts | Specify the number of failed login attempts that are allowed for the web interface and CLI before the account is locked. (1-10, default 0). 0 means that there is no limit. |
| Lockout Time (min) | Specify the number of minutes that a user is locked out (0-60 minutes) if the number of failed attempts is reached. The default 0 means that there is no limit to the number of attempts. |
| **Management Tab** | Configure the management port settings labeled **MGT** on the front of the device. This port is used for all log collector communication. |
| Interface | The interface cannot be changed and the default label is MGT. |
| Speed and Duplex | Select the interface speed in Mbps (10, 100, or 1000) and the interface transmission mode full-duplex (Full), half-duplex (Half), or negotiated automatically (Auto). |
| IP Address | Enter the IP Address of the log collector management interface.<br>The default IP address is 192.168.1.1. |
| Netmask | Enter the network mask for the IP address, such as "255.255.255.0". |
| Default Gateway | Enter the IP address of the default router (must be on the same subnet as the management port). |
| IPv6 Address | Enter the IPv6 address of the log collector management interface. |
| IPv6 Default Gateway | Enter the IPv6 address of the default router (must be on the same subnet as the management port). |
| MTU | Enter the maximum transmission unit (MTU) in bytes for packets sent on this interface (range 512 to 1500, default 1500). |
| Management Interface Services | Select the services that are enabled on the managed interface on the log collector device:<br>• SSH—Select the check box to enable secure shell.<br>• Ping—Select the check box to enable ping.<br>• SNMP—Select the check box to enable the Simple Network Managed Protocol. |

**Table 154 Managed Collectors Page**

| Field | Description |
|---|---|
| Permitted IP Addresses | Click **Add** to enter the list of IP addresses from which management is allowed. |
| **Disks Tab** | Click **Add** to define the RAID 1 disk pair that will be used to store logs. You can then add additional disk pairs as needed to expand your storage capacity.<br>By default, the M-100 is shipped with the first RAID 1 pair enabled with drives installed in bays A1/A2. You can add up to 3 more RAID 1 pairs to increase storage capacity by putting in new RAID 1 pairs in bays B1/B2, C1/C2, and D1/D2. In the software, the RAID 1 pair in bays A1/A2 is named Drive Pair <A\|B\|C\|D>.<br>Refer to "Log Collector Storage" on page 414. |

## Defining Log Collector Groups

▶ *Panorama > Collector Groups*

Collector groups are used to assign Panorama managed firewalls to log collectors that will be used to offload the work of log collection that would normally be handled by the Panorama management server. Once the log collectors are in place and the firewalls are configured, the defined logs for each device will be sent to the log collectors and Panorama will then query the log collectors for aggregated log viewing or investigation. You also use collector groups to define storage retention and SNMP settings.

For detailed descriptions of each of the Collector Groups settings, refer to "Collector Groups Settings" on page 412.

To configure log collector groups:

1. From the Panorama management server, navigate to **Panorama > Collector Groups**.
2. Click **Add** and enter a name to identify the collector group.
3. In the **Min Retention Period (days)** field, enter the number of days that should be maintained across all log collectors in the group before an alert is generated. An alert violation in the form of a system log will be generated if the current date minus the oldest log is less than the defined min retention period.
4. Click the **Monitoring** tab and enter the SNMP settings for your environment in order to manage your devices through a system management solution.
5. Click the **Log Forwarding** tab and in the **Collectors** window, click **Add** and select the log collectors that will be part of this collector group. This list is generated by the collectors defined in **Panorama > Managed Collectors**.

***Note:*** *When you add the device serial number to the collector group, the managed device will start to send all logs to the collector group. To have the managed device revert back to sending logs to the Panorama manager, just remove the device from the collector group. This would also be required when migrating managed devices to a different install of Panorama manager. If you forget to do this, you will need to run the operational command* `delete log-collector preference-list` *from the managed device.*

6. Configure the mapping that defines which devices forward to a preferred list of collectors in preference order. In the **Log Forwarding** tab, click **Add** under **Devices** and **Collectors** list windows and the **Devices** window will appear.
7. In the **Devices** drop-down, select a managed firewall and then under **Collectors** click **Add** and select a log collector. If you select multiple collectors the first collector added will be the primary collector, the second will be the secondary collector, and so on. If the primary fails, logs will be sent to the secondary.
   Click **OK** to save your changes.
8. Click **Commit** and in the **Commit Type** drop down select **Panorama** and click **OK**.
9. Click **Commit** and in the **Commit Type** drop down select **Collector Group** and click **OK**. Once the collector group commit completes, the managed firewalls that have been assigned to collectors will start forwarding logs to the collectors.

   Now that firewall logs are being sent to the log collectors, you then use Panorama's ACC, PDF Reports, and the Logs viewer to query aggregated information for all of your managed firewalls. The ACC will query data directly from the log collectors, which

means that the ACC is looking at forwarded data from the firewall. If you deploy an M-100 as a Panorama manager and log collector, the ACC will also query the data forwarded to the log collector that is on the same device in this case.

**Table 155 Collector Groups Settings**

| Field | Description |
|---|---|
| **General Tab** | |
| Name | Enter a name to identify this collector group name that will be used to group log collectors for configuration and software update purposes (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. |
| Log Storage | Indicates the current log storage quotas for the collector group. If you click on the capacity text, the **Log Storage Settings** window will appear. From here, you can allocate storage to various log features, such as Traffic, Threat, Config, System, and Alarm. You can also click **Restore Defaults** to use the default log allocation settings. |
| Min Retention Period (days) | Specify the retention period in days that should be maintained across all log collectors in the group before an alert is generated. An alert violation in the form of a system log will be generated if the current date minus the oldest log is less than the defined min retention period (range 1-2000 days). |
| **Monitoring Tab** | |

**Table 155 Collector Groups Settings (Continued)**

| Field | Description |
|---|---|
| SNMP | The SNMP option enables you to collect information about the log collectors, including: connection status, Disk drive statistics, software version, average CPU, Average log/second, and storage duration per DB type (e.g. minutes, hours, days, weeks). SNMP information is based on a per-collector group.<br><br>Specify the SNMP settings:<br>• Location—Specify the location of the log collector device.<br>• Contact—Specify an email contact for this device.<br>• Access Setting—Specify the SNMP version that will be used to communicate with the Panorama management server (V2c or V3).<br>If you select V3, specify the following:<br>– Views— Click Add and configure the following settings:<br>› View—Specify a name for a view.<br>› OID—Specify the object identifier (OID).<br>› Option (include or exclude)—Choose whether the OID is to be included or excluded from the view.<br>› Mask—Specify a mask value for a filter on the OID in hexadecimal format (for example, 0xf0)<br>– Users—Click **Add** and configure the following settings:<br>› Users—Specify a user name that will be used for authentication between the log collector and SNMP management server.<br>› View—Specify the group of views for the user.<br>› Authpwd—Specify the user's authentication password (minimum 8 characters). Only Secure Hash Algorithm (SHA) is supported.<br>› Privpwd—Specify the user's encryption password (minimum 8 characters). Only Advanced Encryption Standard (AES) is supported.<br>• SNMP Community—Specify the SNMP community string that is used by your SNMP management environment. SNMPv2c Only (default is public). |
| **Log Forwarding Tab** | |

**Table 155 Collector Groups Settings (Continued)**

| Field | Description |
|---|---|
| Collectors | Click **Add** and select the log collector from the drop-down that will be part of this group. The drop-down will show all log collectors that are available in the **Panorama > Managed Collectors** page. |
| Devices | Click **Add** and then click the **Devices** drop-down and select the managed firewall that will be part of this collector group.<br>Click **Add** in the **Collectors** window and select the collector that you would like to assign this firewall for log forwarding.<br>Click **OK** to save your changes.<br>When viewing the **Devices** window, the **Devices** column will list each firewall and the **Collectors** column will list the assigned collector(s) for the firewall.<br>The first collector you specify will be the primary collector for the firewall. If the primary collector fails, the firewall will then send logs to the secondary collector. If the secondary fails, then the tertiary collector will be used, and so on.<br>**Note:** When you add the device serial number to the collector group, the managed device will start to send all logs to the collector group. To have the managed device revert back to sending logs to the Panorama manager, just remove the device from the collector group. This would also be required when migrating managed devices to a different install of Panorama manager. If you forget to do this, you will need to run the operational command `delete log-collector preference-list` from the managed device. |

### Log Collector Storage

As your environment grows, you will need to increase the storage capacity of the log collector M-100 appliances, or you may need to add new log collectors.

By default, the M-100 is shipped with the first RAID 1 pair enabled with drives installed in bays A1/A2. You can add up to 3 more RAID 1 pairs to increase storage capacity by adding in new RAID 1 pairs in bays B1/B2, C1/C2, and D1/D2. In the software, the RAID 1 pair in bays A1/A2 is named Disk Pair A, B1/B2 is Disk Pair B, C1/C2 is Disk Pair C, D1/D2 is Disk Pair D.

For information on physically replacing drives on the M-100, refer to the *M-100 Hardware Reference Guide*.

***Note:*** *The RAID 1 disk arrays on the M-100 requires that each mirrored disk pair be installed in the appropriate disk bays. The device ships with A1/A2 drives installed and mirrored. To add more disks, you install the next pair into the next Disk Pair, which is bays B1/B2, then C1/C2, and then D1/D2. You cannot install new disks in B1 and C1 for example, and then try to mirror those.*

To expand disk storage on an M-100 using the B1/B2 drive bay:

1. Obtain two identical M-100 disk drives from Palo Alto Networks.
2. You do not have to power down the log collector to add new drives. If you prefer to power down the device, run **`request shutdown system`** from the CLI.

3. Remove the empty drives bays B1/B2 and install the new disks into the drive bays for Disk Pair B.
4. Insert the drive bays back into bays B1/B2.
5. Go to the CLI on the log collector and run **`request system raid add B1`**. This will initialize and format the B1 drive.
6. Run the command again on drive B2 **`request system raid add B2`**. The Raid 1 mirror will be created and the drives are now ready to be added to the collector group from the Panorama management server.

***Note:*** *The time taken to mirror the data on the drive may vary from several minutes to a couple hours, depending on the amount of data on the drive. Use the command* `show system raid detail` *to monitor the progress of the RAID configuration.*

7. Go to the Panorama management server that is managing this log collector, navigate to **Panorama > Manage Collectors** and click the log collector to open the **Collector** window. Select the **Disks** tab.
8. Disk A should already exist. Click **Add** and select **Disk Pair B** and then click **OK** to make the new disk pair available to the system.

When you increase or decrease storage on a log collector, the new storage capacity information is recognized by the Panorama management server and the log collector groups will be updated to re-balance logging. Regardless of firewall to collector assignments, logs are balanced across all collectors in the group to maintain even storage usage among the collector group. Note that during this re-balancing process, you may see heavy network traffic as log data is moved between collectors.

To view the disk capacity for a log collector group, or an individual log collector, go to the web interface on the Panorama management server, navigate to **Panorama > Collector Groups**, click a group name and then click the **General** tab. Next to **Log Storage**, you will see a link that displays the total and free storage information. If you click this link, the **Log Storage Settings** windows will appear where you can allocate storage for various features.

# Viewing Firewall Deployment Information

▶ *Panorama > Device Deployment*

Open the **Device Deployment** pages to view current deployment information on the managed devices and manage software versions on the devices, as described in the following table.

**Table 156. Panorama Deployment Pages**

| Field | Description |
|---|---|
| **Software** | Lists the versions of firewall software that are available for installation on the managed firewalls. |
| **SSL VPN Client** | Lists the versions of SSL VPN client software that are available for installation on the managed firewalls. |
| **GlobalProtect Client** | Lists the versions of GlobalProtect client software that are available for installation on the managed firewalls. |
| **Dynamic Updates** | Lists the threat and application definitions that are available for use on the managed firewalls. Refer to "Updating Threat and Application Definitions" on page 55 for information on using this page. |
| **Licenses** | Lists each managed device and the current license status. Each entry indicates whether the license is active ( icon) or inactive ( icon), along with the expiration date for active licenses.<br>Perform either of the following functions on this page:<br>• Click **Refresh** to update the list.<br>• Click **Activate** to activate a license. Select the managed devices for activation and enter the authentication code that Palo Alto Networks provided for the device. |

Perform any of the following functions on the **Software**, **SSL VPN**, or **GlobalProtect** pages:

- Click **Refresh** to view the latest software releases available from Palo Alto Networks.
- Click **Release Notes** to view a description of the changes in a release.
- Click **Download** to install a new release from the download site. When the download is complete, a checkmark is displayed in the **Downloaded** column. To install a downloaded release, click **Install** next to the release.

  During installation, you are asked whether to reboot when installation is complete. When the installation is complete, you will be logged out while the firewall is restarted. The firewall will be rebooted, if that option was selected.
- Click **Upload** to install or activate a release that you previously stored on your PC. Browse to select the software package, and click **Install from File**. Choose the file that you just selected from the drop-down list, and click **OK** to install the image.
- Click the **Delete** icon to delete an outdated release.

# Backing Up Firewall Configurations

▶ *Panorama > Setup*

Panorama automatically saves every committed configured from the managed firewalls. You can configure the number of versions to keep on the Panorama device by using the Management settings under **Setup** on the **Panorama** tab. The default is 100. For instructions on configuring the number of versions, refer to "Defining Management Settings" on page 30.

To manage backups on Panorama, choose **Panorama > Managed Devices** and click **Manage** in the **Backups** column for a device. A window opens to show the saved and committed configurations for the device. Click a **Load** link to restore the backup to the candidate configuration, and then make any desired changes and click **Commit** to restore the loaded configuration to the device. To remove a saved configuration, click the icon.

# Scheduling Configuration Exports

▶ *Panorama > Scheduled Config Export*

Panorama saves a backup of running configurations from all managed devices in addition to its own running configurations. Use the **Scheduled Config Export** page to collect the running configurations from all of the managed devices, package them in one gzip file, and schedule the package for daily delivery to an FTP server or by using Secure Copy (SCP) to transfer data securely to a remote host. The files are in XML format with file names that are based on the device serial numbers.

Use this page to set up a schedule for collection and export of the managed device configurations.

**Table 157. Scheduling Configuration Bundle Exports**

| Field | Description |
|---|---|
| Name | Enter a name to identify the configuration bundle export job (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, hyphens, and underscores. |
| Description | Enter an optional description. |
| Enable | Select the check box to enable the export job. |
| Scheduled export start time (daily) | Specify the time of day to start the export (24 hour clock, format HH:MM). |
| Protocol | Select the protocol to use to export logs from the firewall to a remote host. You can use SCP to export logs securely, or you can use FTP, which is not a secure protocol. |
| Hostname | Enter the IP address or host name of the target FTP server. |
| Port | Enter the port number on the target server. |
| Path | Specify the path located on the FTP server that will be used to store the exported information. |
| Enable FTP Passive Mode | Select the check box to use FTP passive mode. |
| Username | Specify the user name on the target system. |

**Table 157. Scheduling Configuration Bundle Exports (Continued)**

| Field | Description |
|---|---|
| Password<br>Confirm Password | Specify the password for the user on the target system. |
| Test SCP server connection | Click this button to test communication between the firewall and the SCP host/server.<br>To enable the secure transfer of data, you must verify and accept the host key of the SCP server. The connection is not established until the host key is accepted. |

# Upgrading the Panorama Software

▶ *Panorama > Software*

To upgrade to a new release of Panorama software, you can view the latest versions of the Panorama software available from Palo Alto Networks, read the release notes for each version, and then select the release you want to download and install (a support license is required).

To upgrade the Panorama software, click **Refresh** to view the latest software releases available from Palo Alto Networks. To view a description of the changes in a release, click **Release Notes** next to the release.

***Note:*** *Panorama periodically performs a file system integrity check (FSCK) to prevent corruption of the Panorama system files. This check occurs after 8 reboots or at a reboot that occurs 90 days after the last file system integrity check was executed. If Panorama is running a FSCK, you will see a warning on the web interface and SSH login screens indicating that an FSCK is in progress and you cannot log in until it completes. The time to complete this process varies by the size of the storage system; depending on the size, it can take several hours before you can log back into Panorama.*
*To view progress, set up console access to Panorama.*

1. To install a new release:

   a. Click **Download** next to the release to be installed. When the download is complete, a checkmark is displayed in the **Downloaded** column.

   b. To install a downloaded release, click **Install** next to the release.

   When the installation is complete, you will be logged out while the Panorama system is restarted.

2. To delete an outdated release, click ☒ next to the release.

***Note:*** *Software is deleted to make room for newer version downloads. This happens automatically and cannot be manually controlled.*

Chapter 14

# Configuring WildFire

This chapter describes how to use WildFire for analysis and reporting on malware that traverses the firewall:


- "About WildFire" in the next section
- "Setting Up WildFire on the Firewall" on page 423
- "Configuring WildFire Forwarding" on page 424
- "WildFire Data Filtering Log" on page 425
- "Using the WildFire Portal" on page 425


## About WildFire

WildFire allows you to securely send files to the Palo Alto Networks cloud-based malware analysis center, where the files are automatically analyzed in a virtual sandbox for malicious activity. Files can be submitted automatically from firewalls through the use of a file blocking policy, manually through the WildFire Portal, or using the WildFire API (WildFire subscription required). The WildFire system executes the file in a virtual environment and watches for many potentially malicious behaviors and techniques, such as modifying critical system files, disabling security features, or using a variety of methods to evade detection. As new malware is detected, WildFire automatically generates antivirus signatures and distributes these signatures to firewalls with a threat prevention subscription via the daily antivirus signature updates. If you are a WildFire subscriber, signatures are available on a sub-hourly basis.

Supported file types include Win32 Portable Executable (PE) files (e.g. exe, dll, and scr). When choosing file types in the objects profile, you can choose **PE** to cover all Win32 PE file types.

File types can be analyzed even if they are compressed (zip, gzip) or over SSL if decryption is enabled in the policy. You do not need to add (zip, gzip) to the file blocking profile to support analysis of compressed PE files since this is done automatically.

***Note:*** *When choosing **PE** in the objects profile **File Types** column to select a category of file types, do not add an individual file type that is part of that category or you will receive redundant entries in the **Data Filtering** logs. For example, if you choose **PE** you would not want to additionally include **exe**. This also applies to the **zip** file type, since supported file types that are zipped will automatically be sent to WildFire.*

*Choosing a category will also ensure that as new file type support is added to a given category, they will automatically be part of your object profile. You can also select **any** to cause all supported file types to be uploaded to WildFire.*

As new malware is detected, WildFire will automatically generate antivirus signatures and will distribute these signatures to firewalls that have a threat prevention subscription. The signatures will be part of the daily antivirus signatures that are updated every 24-48 hours.

If you have a WildFire subscription, you will have several new features enabled, including:

- **WildFire Dynamic Updates**—A new WildFire section in **Device > Dynamic Updates** is now available. As new malware is detected, WildFire signatures are generated every 30 minutes, and the firewall's WildFire update schedule can be configured to poll for new WildFire signatures every 15, 30, or 60 minutes. As new WildFire signatures are downloaded to your firewall, you can configure the firewall to take specific actions on these signatures, separate from the regular antivirus signature actions in the antivirus profile.

- **Integrated WildFire Logs**—As files are uploaded and analyzed by WildFire, log data is sent back to the device after analysis, along with the analysis results. Logs are written to **Monitor > Logs > WildFire**.

- **WildFire API**—The WildFire license provides access to the WildFire API, which allows for programmatic access to the WildFire service in the cloud, apart from use by Palo Alto Networks firewalls. The WildFire API can be used to submit, analyze, and review reports for files submitted to the WildFire system. You can upload up to 100 files per day and query the WildFire portal up to 1000 times per day.

Results of the detailed analysis of the submitted files are also available through the WildFire portal and you can view this information without a subscription. You can use the WildFire portal to see which users were targeted, the applications that were used, and the malicious behavior that was observed. You can also configure the WildFire portal to send email notifications when results are available for review. Refer to “Using the WildFire Portal” on page 425.

# Setting Up WildFire on the Firewall

Perform these tasks to set up your environment to use WildFire.

1. On the firewall, configure WildFire settings on the **Device > Setup** page. Refer to "Configuring WildFire Settings on the Firewall" on page 423.
2. On the firewall, configure your file blocking profiles to include the **forward** or **continue-and-forward** action. Refer to "Configuring WildFire Forwarding" on page 424 for procedures, or refer to "File Blocking Profiles" on page 220 for more information.
3. Incorporate the file blocking profiles in a security policy, as you would for any other file blocking profiles. Refer to "Security Policies" on page 187.
4. Access the WildFire portal and configure optional settings. Refer to "Using the WildFire Portal" on page 425.

You can now access the WildFire portal to view reports. Refer to "Viewing WildFire Reports" on page 427.

# Configuring WildFire Settings on the Firewall

▶ *Device > Setup > WildFire*

Use the **WildFire** tab to control the information to be sent to the WildFire server.

***Note:*** *To forward decrypted content to WildFire, you need to select the "Allow Forwarding of Decrypted Content" check box in* ***Device > Setup > Content-ID > URL Filtering*** *Settings box.*

**Table 158. WildFire Settings on the Firewall**

| Field | Description |
|---|---|
| **General Settings** | |
| WildFire Server | Specify the URL of a WildFire server. Specify the value **default-cloud** to allow the firewall to automatically find the closest WildFire server. |
| Maximum File Size (MB) | Specify the maximum file size that will be forwarded to the WildFire server (range 1-10 MB, default 2 MB). Files larger than the specified size are not sent. |

**Table 158. WildFire Settings on the Firewall (Continued)**

| Field | Description |
|---|---|
| **Session Information Settings** | |
| Settings | Specify the information to be forwarded to the WildFire server. By default, all are selected:<br>• **Source IP**—Source IP address that sent the suspected file.<br>• **Source Port**—Source port that sent the suspected file.<br>• **Destination IP**—Destination IP address for the suspected file.<br>• **Destination Port**—Destination port for the suspected file.<br>• **Vsys**—Firewall virtual system that identified the possible malware.<br>• **Application**—User application that was used to transmit the file.<br>• **User**—Targeted user.<br>• **URL**—URL associated with the suspected file.<br>• **Filename**—Name of the file that was sent. |

# Configuring WildFire Forwarding

After configuring the WildFire settings on the firewall, you are now ready to configure forwarding to allow files to be sent to the WildFire system for analysis. You first create a file blocking profile and then include that profile in a security policy.

To configure a file blocking profile for WildFire:

1. Navigate to **Objects > Security Profiles > File Blocking**.
2. Click **Add** to add a new profile and enter a **Name** and **Description**.
3. Click **Add** in the **File Blocking Profile** window. Click in the **Names** field and enter a rule name.
4. Select the **Application(s)** that will be identified for this profile. For example, if you choose the application **web-browsing**, the profile will identify files downloaded when the user downloads files from a web page.
5. In the **File Type** field select the file types that you would like to forward for analysis.
6. In the **Direction** field, select **upload**, **download**, or **both**. If you select **both,** this will identify files that are uploaded or downloaded by the user.
7. In the **Action** field, select **forward**. This will cause any identified file to be forwarded to the WildFire system, analysis will occur and the file will then be delivered to the user. If you select **continue-and-forward**, the user will be prompted with a continue page before the download will occur.

If WildFire determines that the file is a malware file, a report will be generated on the WildFire Portal and a new antivirus signature will be created within 24-48 hours. If you are a WildFire subscriber, you will also see a **WildFire** log entry on the firewall and a signature will be created within an hour.

To configure a security policy for WildFire:

1. Navigate to **Policies** > **Security.**

2. Click on **Add** to create a new policy, or select an existing policy.
3. Click the **Actions** tab and under **Profile Setting** click the drop-down next to **File Blocking** and choose the Security Profile that you created.
4. Click **Commit** to active your changes.

***Note:*** *When you create a file blocking profile with the action* ***continue*** *or* ***continue-and-forward*** *(used for WildFire forwarding), you can only choose the application* ***web-browsing****. If you choose any other application, traffic that matches the security policy will not flow through the firewall due to the fact that the users will not be prompted with a continue page.*

# WildFire Data Filtering Log

This section describes the various forwarding actions that you will see in the logs when WildFire is configured and files are forwarded for analysis. The logs will be located in Monitor > Logs > Data Filtering.

**Table 159. WildFire Data Filtering Log Descriptions**

| Log | Description |
|---|---|
| wildfire-upload-success | The file was sent to the cloud. This means the file is not signed by a trusted file signer, it has not been previously analyzed by WildFire. |
| wildfire-upload-skip | This action will be displayed for all files identified as eligible to be sent to WildFire by a file blocking profile/security policy, but did not need to be analyzed by WildFire because it has already been analyzed previously. In this case, the "forward" action will appear in the Data Filtering log because it was a valid forward action, but it was not sent to WildFire and analyzed because the file has already been sent to the WildFire cloud from another session, possibly from another firewall. |

# Using the WildFire Portal

To access the WildFire portal, go to *https://wildfire.paloaltonetworks.com* and log in using your Palo Alto Networks support credentials or your WildFire account.

The portal opens to display the dashboard, which lists summary report information for all of the firewalls associated with the specific WildFire license or support account (as well as any files that have been uploaded manually). The display includes the number of analyzed files and indicates how many are infected with malware, are benign, or are pending analysis.

***Note:*** *To upload a file manually, click* ***Upload File*** *in the upper right corner of the WildFire page.*

**Figure 55. WildFire Dashboard**

# Configuring Settings on the WildFire Portal

Click the **Settings** link at the top of the WildFire portal to configure the following settings:

**Table 160. Settings on the WildFire Portal**

| Field | Description |
|---|---|
| Password | To change your password, enter values in the following fields:<br>• **Current Password**—Enter your current password.<br>• **New Password/Confirm Password**—Enter and then reenter a new password. |
| Time Zone | Select the time zone from the drop-down list. This is the time zone that is used to indicate when WildFire receives files. |

**Table 160. Settings on the WildFire Portal (Continued)**

| Field | Description |
|---|---|
| Email Notifications | Choose the email notifications that you would like to receive. The email notifications are sent to the currently logged in WildFire user. For each device, and for files that are manually uploaded to the WildFire server, you can choose any of the following email notifications:<br>• **Malware**—Select this check box to have an Email notification sent when the file is determined to be malware.<br>• **Benign**—Select this check box to have an Email notification sent when the file is determined to be benign.<br>***Note:*** *You can also select both check boxes to receive Email notifications for all files (malware/benign).* |

# Viewing WildFire Reports

Click the **Reports** button at the top of the WildFire portal to view the list of available reports. Search options are available at the top of the page, and pagination controls are included.

To view an individual report, click the icon to the left of the report name. To print a detailed report, use the print option on your browser.

**Figure 56. WildFire Reports Page**

# Appendix A
# Custom Pages

Custom response pages allow you to notify end users of policy violations and special access conditions. Each page can include references to the user's IP address, the URL for which access is attempted, and the URL category. These parameters can also be used in links to trouble-ticketing systems.

This appendix provides HTML code for the following default custom response pages:


- "Default Antivirus Response Page" in the next section
- "Default Application Block Page" on page 431
- "Default File Blocking Block Page" on page 431
- "Default URL Filtering Response Page" on page 432
- "Default Anti-spyware Download Response Page" on page 433
- "Default Decryption Opt-out Response Page" on page 433
- "Captive Portal Comfort Page" on page 434
- "URL Filtering Continue and Override Page" on page 434
- "SSL VPN Login Page" on page 435
- "SSL Certificate Revoked Notify Page" on page 436

***Note:*** *For information on importing and exporting custom response pages, refer to "Defining Custom Response Pages" on page 115.*

## Default Antivirus Response Page

```
<html>

<head>
<meta http-equiv=Content-Type content="text/html; charset=windows-1252">
<meta name=Generator content="Microsoft Word 11 (filtered)">
<title>This is a test</title>
<style>
<!--
 /* Font Definitions */
 @font-face
    {font-family:"Microsoft Sans Serif";
```

```
    panose-1:2 11 6 4 2 2 2 2 2 4;}
 /* Style Definitions */
 p.MsoNormal, li.MsoNormal, div.MsoNormal
    {margin:0in;
    margin-bottom:.0001pt;
    font-size:12.0pt;
    font-family:"Times New Roman";}
h4
    {margin-top:12.0pt;
    margin-right:0in;
    margin-bottom:3.0pt;
    margin-left:0in;
    page-break-after:avoid;
    font-size:14.0pt;
    font-family:"Times New Roman";}
p.SanSerifName, li.SanSerifName, div.SanSerifName
    {margin:0in;
    margin-bottom:.0001pt;
    text-autospace:none;
    font-size:10.0pt;
    font-family:"Microsoft Sans Serif";
    font-weight:bold;}
p.BoldNormal, li.BoldNormal, div.BoldNormal
    {margin:0in;
    margin-bottom:.0001pt;
    font-size:12.0pt;
    font-family:"Times New Roman";
    font-weight:bold;}
span.Heading10
    {color:black
    font-weight:bold;}
p.SubHeading1, li.SubHeading1, div.SubHeading1
    {margin-top:12.0pt;
    margin-right:0in;
    margin-bottom:3.0pt;
    margin-left:0in;
    page-break-after:avoid;
    font-size:12.0pt;
    font-family:"Times New Roman";
    font-weight:bold;}
@page Section1
    {size:8.5in 11.0in;
    margin:1.0in 1.25in 1.0in 1.25in;}
div.Section1
    {page:Section1;}
-->
</style>

</head>

<body lang=EN-US>

<div class=Section1>

<p class=MsoNormal>This is a test.</p>

</div>

</body>

</html>
```

## Default Application Block Page

```
<html>
<head>
<title>Application Blocked</title>
<style>
#content{border:3px solid#aaa;background-
color:#fff;margin:40;padding:40;font-family:Tahoma,Helvetica,Arial,sans-
serif;font-size:12px;}
  h1{font-size:20px;font-weight:bold;color:#196390;}
  b{font-weight:bold;color:#196390;}
</style>
</head>
<body bgcolor="#e7e8e9">
<div id="content">
<h1>Application Blocked</h1>
<p>Access to the application you were trying to use has been blocked in
accordance with company policy. Please contact your system administrator if
you believe this is in error.</p>
<p><b>User:</b> <user/> </p>
<p><b>Application:</b> <appname/> </p>
</div>
</body>
</html>
```

## Default File Blocking Block Page

```
<html>

<head>
<meta http-equiv=Content-Type content="text/html; charset=windows-1252">
<meta name=Generator content="Microsoft Word 11 (filtered)">
<title>This is a test</title>
<style>
<!--
 /* Font Definitions */
 @font-face
     {font-family:"Microsoft Sans Serif";
     panose-1:2 11 6 4 2 2 2 2 2 4;}
 /* Style Definitions */
 p.MsoNormal, li.MsoNormal, div.MsoNormal
     {margin:0in;
     margin-bottom:.0001pt;
     font-size:12.0pt;
     font-family:"Times New Roman";}
h4
     {margin-top:12.0pt;
     margin-right:0in;
     margin-bottom:3.0pt;
     margin-left:0in;
     page-break-after:avoid;
     font-size:14.0pt;
     font-family:"Times New Roman";}
p.SanSerifName, li.SanSerifName, div.SanSerifName
     {margin:0in;
     margin-bottom:.0001pt;
     text-autospace:none;
     font-size:10.0pt;
     font-family:"Microsoft Sans Serif";
     font-weight:bold;}
p.BoldNormal, li.BoldNormal, div.BoldNormal
     {margin:0in;
     margin-bottom:.0001pt;
     font-size:12.0pt;
     font-family:"Times New Roman";
     font-weight:bold;}
```

```
span.Heading10
     {color:black
     font-weight:bold;}
p.SubHeading1, li.SubHeading1, div.SubHeading1
     {margin-top:12.0pt;
     margin-right:0in;
     margin-bottom:3.0pt;
     margin-left:0in;
     page-break-after:avoid;
     font-size:12.0pt;
     font-family:"Times New Roman";
     font-weight:bold;}
@page Section1
     {size:8.5in 11.0in;
     margin:1.0in 1.25in 1.0in 1.25in;}
div.Section1
     {page:Section1;}
-->
</style>

</head>

<body lang=EN-US>

<div class=Section1>

<p class=MsoNormal>This is a test.</p>

</div>

</body>

</html>
```

# Default URL Filtering Response Page

```
<html>
<head>
<title>Web Page Blocked</title>
<style>
#content{border:3px solid#aaa;background-
color:#fff;margin:40;padding:40;font-family:Tahoma,Helvetica,Arial,sans-
serif;font-size:12px;}
  h1{font-size:20px;font-weight:bold;color:#196390;}
  b{font-weight:bold;color:#196390;}
</style>
</head>
<body bgcolor="#e7e8e9">
<div id="content">
<h1>Web Page Blocked</h1>
<p>Access to the web page you were trying to visit has been blocked in
accordance with company policy. Please contact your system administrator if
you believe this is in error.</p>
<p><b>User:</b> <user/> </p>
<p><b>URL:</b> <url/> </p>
<p><b>Category:</b> <category/> </p>
</div>
</body>
</html>
```

## Default Anti-spyware Download Response Page

```
<application-type>
    <category>
        <entry name="networking" id="1">
            <subcategory>
                <entry name="remote-access" id="1"/>
                <entry name="proxy" id="2"/>
                <entry name="encrypted-tunnel" id="3"/>
                <entry name="routing" id="4"/>
                <entry name="infrastructure" id="5"/>
                <entry name="ip-protocol" id="6"/>
            </subcategory>
        </entry>
        <entry name="collaboration" id="2">
            <subcategory>
                <entry name="email" id="7"/>
                <entry name="instant-messaging" id="8"/>
                <entry name="social-networking" id="9"/>
                <entry name="internet-conferencing" id="10"/>
                <entry name="voip-video" id="11"/>
            </subcategory>
        </entry>
        <entry name="media" id="3">
            <subcategory>
                <entry name="video" id="12"/>
                <entry name="gaming" id="13"/>
                <entry name="audio-streaming" id="14"/>
            </subcategory>
        </entry>
        <entry name="business-systems" id="4">
            <subcategory>
                <entry name="auth-service" id="15"/>
                <entry name="database"id="16"/>
                <entry name="erp-crm" id="17"/>
                <entry name="general-business" id="18"/>
                <entry name="management" id="19"/>
                <entry name="office-programs" id="20"/>
                <entry name="software-update" id="21"/>
                <entry name="storage-backup" id="22"/>
            </subcategory>
         </entry>
         <entry name="general-internet" id="5">
            <subcategory>
                <entry name="file-sharing" id="23"/>
                <entry name="internet-utility" id="24"/>
            </subcategory>
        </entry>
    </category>
    <technology>
          <entry name="network-protocol" id="1"/>
          <entry name="client-server" id="2"/>
          <entry name="peer-to-peer" id="3"/>
          <entry name="web-browser" id="4"/>
    </technology>
</application-type>
```

## Default Decryption Opt-out Response Page

```
<h1>SSL Inspection</h1>
<p>In accordance with company security policy, the SSL encrypted connection
you have initiated will be temporarily unencrypted so that it can be
inspected for viruses, spyware, and other malware.</p>
<p>After the connection is inspected it will be re-encrypted and sent to its
destination. No data will be stored or made available for other purposes.</p>
<p><b>IP:</b> <url/> </p>
<p><b>Category:</b> <category/> </p>
```

## Captive Portal Comfort Page

```
<h1 ALIGN=CENTER>Captive Portal</h1>

<h2 ALIGN=LEFT>In accordance with company security policy, you have to
authenticate before accessing the network.</h2>

<pan_form/>
```

## URL Filtering Continue and Override Page

```
<html>
<head>
<title>Web Page Blocked</title>
<style>
#content{border:3px solid#aaa;background-
color:#fff;margin:40;padding:40;font-family:Tahoma,Helvetica,Arial,sans-
serif;font-size:12px;}
  h1{font-size:20px;font-weight:bold;color:#196390;}
  b{font-weight:bold;color:#196390;}
        form td, form input {
                font-size: 11px;
                font-weight: bold;
        }
        #formtable {
                height: 100%;
                width: 100%;
        }
        #formtd {
                vertical-align: middle;
        }
        #formdiv {
                margin-left: auto;
                margin-right: auto;
        }
</style>
<script type="text/javascript">
function pwdCheck() {
    if(document.getElementById("pwd")) {
         document.getElementById("continueText").innerHTML = "If you require
access to this page, have an administrator enter the override password
here:";
    }
}
</script>
</head>
<body bgcolor="#e7e8e9">
<div id="content">
<h1>Web Page Blocked</h1>
<p>Access to the web page you were trying to visit has been blocked in
accordance with company policy. Please contact your system administrator if
you believe this is in error.</p>
<p><b>User:</b> <user/> </p>
<p><b>URL:</b> <url/> </p>
<p><b>Category:</b> <category/> </p>

<hr>
<p id="continueText">If you feel this page has been incorrectly blocked, you
may click Continue to proceed to the page. However, this action will be
logged.</p>
<div id="formdiv">
<pan_form/>
</div>
<a href="#" onclick="history.back();return false;">Return to previous page</
a>
</div>
</body>
</html>
```

# SSL VPN Login Page

```
<HTML>
<HEAD>
<TITLE>Palo Alto Networks - SSL VPN</TITLE>
<meta http-equiv="Content-Type" content="text/html; charset=iso-8859-1">
<link rel="stylesheet" type="text/css" href="/styles/
falcon_content.css?v=@@version">
<style>
td {
     font-family: Verdana, Arial, Helvetica, sans-serif;
     font-weight: bold;
     color: black; /*#FFFFFF; */
}
.msg {
    background-color: #ffff99;
    border-width: 2px;
    border-color: #ff0000;
    border-style: solid;
    padding-left: 20px;
    padding-right: 20px;
    max-height: 150px;
    height: expression( this.scrollHeight > 150 ? "150px" : "auto" ); /* sets
max-height for IE */
    overflow: auto;
}
.alert {font-weight: bold;color: red;}

</style>
</HEAD>
<BODY bgcolor="#F2F6FA">
     <table style="background-color: white; width:100%; height:45px; border-
bottom: 2px solid #888888;">
          <tr style="background-image:url(/images/logo_pan_158.gif);
background-repeat: no-repeat">
               <td align="left"> </td>
          </tr>
     </table>

     <div align="center">
          <h1>Palo Alto Networks - SSL VPN Portal</h1>
     </div>

<div id="formdiv">
<pan_form/>
</div>
</BODY>
</HTML>
```

# SSL Certificate Revoked Notify Page

```
<META HTTP-EQUIV="Content-Type" CONTENT="text/html; charset=iso-8859-1">
<html>

<head>

<title>Certificate Error</title>

<style>


#content{border:3px solid#aaa;background-
color:#fff;margin:40;padding:40;font-family:Tahoma,Helvetica,Arial,sans-
serif;font-size:12px;}


  h1{font-size:20px;font-weight:bold;color:#196390;}


  b{font-weight:bold;color:#196390;}


</style>
</head>

<body bgcolor="#e7e8e9">

<div id="content">

<h1>Certificate Error</h1>

<p>There is an issue with the SSL certificate of the server you are trying to
contact.</p>

<p><b>Certificate Name:</b> <certname/> </p>

<p><b>IP:</b> <url/> </p>

<p><b>Issuer:</b> <issuer/> </p>

<p><b>Status:</b> <status/> </p>

<p><b>Reason:</b> <reason/> </p>

</div>

</body>

</html>
```

# Appendix B
# Application Categories, Subcategories, Technologies, and Characteristics

The appendix lists application-related categories defined by Palo Alto Networks:

- "Application Categories and Subcategories" in the next section
- "Application Technologies" on page 439
- "Application Characteristics" on page 439

The Applipedia database can also be found on the firewall in **Objects > Applications**, online at http://apps.paloaltonetworks.com/applipedia/ and an App is available on the Apple App Store.

## Application Categories and Subcategories

The following application categories and subcategories are supported:

- business-system
  - auth-service
  - database
  - erp-crm
  - general-business
  - management
  - office-programs
  - software-update
  - storage-backup
- collaboration
  - email
  - instant-messaging
  - internet-conferencing

  - social-business
  - internet-utility
  - social-networking
  - voip-video
  - web-posting
- general-internet
  - file-sharing
  - internet-utility
- media
  - audio-streaming
  - gaming
  - photo-video
- networking
  - encrypted-tunnel
  - infrastructure
  - ip-protocol
  - proxy
  - remote-access
  - routing
- unknown

# Application Technologies

The following application technologies are supported.

**Table 161. Application Technologies**

| Item | Description |
|---|---|
| browser-based | An application that relies on a web browser to function. |
| client-server | An application that uses a client-server model where one or more clients communicate with a server in the network. |
| network-protocol | An application that is generally used for system to system communication that facilitates network operation. This includes most of the IP protocols. |
| peer-to-peer | An application that communicates directly with other clients to transfer information instead of relying on a central server to facilitate the communication. |

# Application Characteristics

The following application characteristics are supported.

**Table 162. Application Characteristics**

| Item | Description |
|---|---|
| Evasive | Uses a port or protocol for something other than its originally intended purpose with the hope that it will traverse a firewall. |
| Excessive Bandwidth | Consumes at least 1 Mbps on a regular basis through normal use. |
| Prone to Misuse | Often used for nefarious purposes or is easily set up to expose more than the user intended. |
| Transfers Files | Has the capability to transfer a file from one system to another over a network. |
| Tunnels Other Apps | Is able to transport other applications inside its protocol. |
| Used by Malware | Malware has been known to use the application for propagation, attack, or data theft, or is distributed with malware. |
| Vulnerability | Has publicly reported vulnerabilities. |
| Widely Used | Likely has more than 1,000,000 users. |
| Continue Scanning for Other Applications | Instructs the firewall to continue looking to see if other application signatures match. If this option is not selected, the first matching signature is reported and the firewall stops looking for additional matching applications. |

# Appendix C
# Common Criteria/Federal Information Processing Standards Support

You can configure the firewall to support the Common Criteria Evaluation Assurance Level 4+ (CCEAL4+) and the Federal Information Processing Standards 140-2 (FIPS 140-2), which are security certifications that ensure a standard set of security assurances and functionalities. These certifications are often required by civilian U.S. government agencies and government contractors.

For more details on these certifications, refer to the following documents:

- **FIPS**—http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140sp/140sp1877.pdf
- **CCEAL4+**—http://www.niap-ccevs.org/st/vid10392/

***Note:*** *PAN-OS 5.0 certification is in process. Visit http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140InProcess.pdf and search for Palo Alto Networks in the Cryptographic Module Validation Program FIPS 140-1 and FIPS 140-2 Modules In Process List for details.*

## Enabling CC/FIPS Mode

Use the following procedure to enable CC/FIPS mode on a software version that supports CC/FIPS. Keep in mind that when you enable CC/FIPS, the device will be reset the factory default settings; all configuration will be removed.

1. Boot the firewall into maintenance mode as follows:
   a. Establish a serial connection with the firewall.
   b. Reboot the device and press the **m** key on the keyboard when you see the following prompt: **Press m to boot to maint partition**.
   c. Press any key on your keyboard when prompted to stop the automatic boot, and then select **PANOS (maint)** as the booting partition.

      For more detailed information about maintenance mode, refer to the *PAN-OS Command Line Interface Reference Guide*.
2. Select **Set CCEAL4 Mode** from the menu.
3. Select **Enable CCEAL4 Mode** from the menu.

4. When prompted, select **Reboot**.

   After successfully switching to CC/FIPS mode, the following status displays: **`CCEAL4 mode enabled successfully`**. In addition, **`CC`** will display at all times in the status bar at the bottom of the web interface. In addition, the console port will now be available as a status output port only. In addition, the default admin login credentials change to admin/paloalto.

# CC/FIPS Security Functions

When CC/FIPS is enabled, the following apply:

- To log into the firewall, the browser must be TLS 1.0 compatible.
- All passwords on the firewall must be at least six characters.
- Accounts are locked after the number of failed attempts that is configured on the **Device > Setup > Management** page. If the firewall is not in CC/FIPS mode, it can be configured so that it never locks out; however in CC/FIPS mode, and lockout time is required.
- The firewall automatically determines the appropriate level of self-testing and enforces the appropriate level of strength in encryption algorithms and cipher suites.
- Non-CC/FIPS approved algorithms are not decrypted and are thus ignored during decryption.
- When configuring IPSec, a subset of the normally available cipher suites is available.
- Self-generated and imported certificates must contain public keys that are 2048 bits (or more).
- The serial port is disabled.
- Telnet, TFTP, and HTTP management connections are unavailable.
- Surf control is not supported.
- High availability (HA) encryption is required.
- PAP authentication is disabled.

# Appendix D
# Open Source Licenses

The software included in this product contains copyrighted software that is licensed under the General Public License (GPL). A copy of that license is included in this document. You may obtain the complete Corresponding Source code from us for a period of three years after our last shipment of this product by sending a money order or check for $5 to:

Palo Alto Networks
Open Source Request
4401 Great America Parkway
Santa Clara, Ca. 95054

Some components of this product may be covered under one or more of the open source licenses listed in this appendix:


- "Artistic License" on page 446
- "BSD" on page 447
- "GNU General Public License" on page 448
- "GNU Lesser General Public License" on page 452
- "MIT/X11" on page 458
- "OpenSSH" on page 458
- "PSF" on page 462
- "PHP" on page 462
- "Zlib" on page 463

# Artistic License

This document is freely plagiarized from the 'Artistic License', distributed as part of the Perl v4.0 kit by Larry Wall, which is available from most major archive sites

This documents purpose is to state the conditions under which these Packages (See definition below) viz: "Crack", the Unix Password Cracker, and "CrackLib", the Unix Password Checking library, which are held in copyright by Alec David Edward Muffett, may be copied, such that the copyright holder maintains some semblance of artistic control over the development of the packages, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

Definitions:

A "Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification, or segments thereof.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

“Freely Available” means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.

2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.

3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when AND WHY you changed that file, and provided that you do at least ONE of the following:

a) place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as uunet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.

b) use the modified Package only within your corporation or organization.

c) rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide separate documentation for each non-standard executable that clearly documents how it differs from the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:

a) distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.

b) accompany the distribution with the machine-readable source of the Package with your modifications.

c) accompany any non-standard executables with their corresponding Standard Version executables, giving the non-standard executables non-standard names, and clearly documenting the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. YOU MAY NOT CHARGE A FEE FOR THIS PACKAGE ITSELF. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that YOU DO NOT ADVERTISE this package as a product of your own.

6. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

7. THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.

# BSD

The following copyright holders provide software under the BSD license:

- Julian Steward
- Thai Open Source Software Center Ltd
- The Regents of the University of California
- Nick Mathewson
- Niels Provos
- Dug Song
- Todd C. Miller
- University of Cambridge
- Sony Computer Science Laboratories Inc.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

# GNU General Public License

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble:

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

# GNU Lesser General Public License

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble:

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

* a) The modified work must itself be a software library.

* b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.

* c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.

* d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

* a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

* b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

* c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

* d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

* e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

* a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

* b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license

would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A

FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

# MIT/X11

Copyright (C) 2001-2002 Daniel Veillard. All Rights Reserved.

Copyright (C) 2001-2002 Thomas Broyer, Charlie Bozeman and Daniel Veillard. All Rights Reserved.

Copyright (C) 1998 Bjorn Reese and Daniel Stenberg.

Copyright (C) 2000 Gary Pennington and Daniel Veillard.

Copyright (C) 2001 Bjorn Reese <breese@users.sourceforge.net>

Copyright (c) 2001, 2002, 2003 Python Software Foundation

Copyright (c) 2004-2008 Paramjit Oberoi <param.cs.wisc.edu>

Copyright (c) 2007 Tim Lauridsen <tla@rasmil.dk>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

# OpenSSH

This file is part of the OpenSSH software.

The licences which components of this software fall under are as follows. First, we will summarize and say that all components are under a BSD licence, or a licence more free than that.

OpenSSH contains no GPL code.

1) Copyright (c) 1995 Tatu Ylonen <ylo@cs.hut.fi>, Espoo, Finland

All rights reserved

As far as I am concerned, the code I have written for this software can be used freely for any purpose. Any derived versions of this software must be clearly marked as such, and if the derived work is incompatible with the protocol description in the RFC file, it must be called by a name other than "ssh" or "Secure Shell".

[Tatu continues]

However, I am not implying to give any licenses to any patents or copyrights held by third parties, and the software includes parts that are not under my direct control. As far as I know, all included source code is used in accordance with the relevant license agreements and can be used freely for any purpose (the GNU license being the most restrictive); see below for details.

[However, none of that term is relevant at this point in time. All of these restrictively licenced software components which he talks about have been removed from OpenSSH, i.e.,

-RSA is no longer included, found in the OpenSSL library

-IDEA is no longer included, its use is deprecated

-DES is now external, in the OpenSSL library

-GMP is no longer used, and instead we call BN code from OpenSSL

-Zlib is now external, in a library

-The make-ssh-known-hosts script is no longer included

-TSS has been removed

-MD5 is now external, in the OpenSSL library

-RC4 support has been replaced with ARC4 support from OpenSSL

-Blowfish is now external, in the OpenSSL library

[The licence continues]

Note that any information and cryptographic algorithms used in this software are publicly available on the Internet and at any major bookstore, scientific library, and patent office worldwide. More information can be found e.g. at "http://www.cs.hut.fi/crypto".

The legal status of this program is some combination of all these permissions and restrictions. Use only at your own responsibility. You will be responsible for any legal consequences yourself; I am not making any claims whether possessing or using this is legal or not in your country, and I am not taking any responsibility on your behalf.

NO WARRANTY

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,

REPAIR OR CORRECTION.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

2) The 32-bit CRC compensation attack detector in deattack.c was contributed by CORE SDI S.A. under a BSD-style license.

Cryptographic attack detector for ssh - source code

Copyright (c) 1998 CORE SDI S.A., Buenos Aires, Argentina.

All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that this copyright notice is retained.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES ARE DISCLAIMED. IN NO EVENT SHALL CORE SDI S.A. BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES RESULTING FROM THE USE OR MISUSE OF THIS SOFTWARE.

Ariel Futoransky <futo@core-sdi.com> <http://www.core-sdi.com>

3) ssh-keyscan was contributed by David Mazieres under a BSD-style license.

Copyright 1995, 1996 by David Mazieres <dm@lcs.mit.edu>.

Modification and redistribution in source and binary forms is permitted provided that due credit is given to the author and the OpenBSD project by leaving this copyright notice intact.

4) The Rijndael implementation by Vincent Rijmen, Antoon Bosselaers and Paulo Barreto is in the public domain and distributed with the following license:

@version 3.0 (December 2000)

Optimised ANSI C code for the Rijndael cipher (now AES)

@author Vincent Rijmen <vincent.rijmen@esat.kuleuven.ac.be>

@author Antoon Bosselaers <antoon.bosselaers@esat.kuleuven.ac.be>

@author Paulo Barreto <paulo.barreto@terra.com.br>

This code is hereby placed in the public domain.

THIS SOFTWARE IS PROVIDED BY THE AUTHORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

5) One component of the ssh source code is under a 3-clause BSD license, held by the University of California, since we pulled these parts from original Berkeley code.

Copyright (c) 1983, 1990, 1992, 1993, 1995

The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS

BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

6) Remaining components of the software are provided under a standard 2-term BSD licence with the following names as copyright holders:

-Markus Friedl

-Theo de Raadt

-Niels Provos

-Dug Song

-Aaron Campbell

-Damien Miller

-Kevin Steves

-Daniel Kouril

-Wesley Griffin

-Per Allansson

-Nils Nordman

-Simon Wilkinson

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

# PSF

1. This LICENSE AGREEMENT is between the Python Software Foundation ("PSF"), and the Individual or Organization ("Licensee") accessing and otherwise using Python 2.3 software in source or binary form and its associated documentation.

2. Subject to the terms and conditions of this License Agreement, PSF hereby grants Licensee a nonexclusive, royalty-free, world-wide license to reproduce, analyze, test, perform and/or display publicly, prepare derivative works, distribute, and otherwise use Python 2.3 alone or in any derivative version, provided, however, that PSF's License Agreement and PSF's notice of copyright, i.e., "Copyright (c) 2001, 2002, 2003 Python Software Foundation; All Rights Reserved" are retained in Python 2.3 alone or in any derivative version prepared by Licensee.

3. In the event Licensee prepares a derivative work that is based on or incorporates Python 2.3 or any part thereof, and wants to make the derivative work available to others as provided herein, then Licensee hereby agrees to include in any such work a brief summary of the changes made to Python 2.3.

4. PSF is making Python 2.3 available to Licensee on an "AS IS" basis. PSF MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, PSF MAKES NO AND DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF PYTHON 2.3 WILL NOT INFRINGE ANY THIRD PARTY RIGHTS.

5. PSF SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON 2.3 FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS A RESULT OF MODIFYING, DISTRIBUTING, OR OTHERWISE USING PYTHON 2.3, OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.

6. This License Agreement will automatically terminate upon a material breach of its terms and conditions.

7. Nothing in this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between PSF and Licensee. This License Agreement does not grant permission to use PSF trademarks or trade name in a trademark sense to endorse or promote products or services of Licensee, or any third party.

8. By copying, installing or otherwise using Python 2.3, Licensee agrees to be bound by the terms and conditions of this License Agreement.

# PHP

The PHP License, version 3.01

Copyright (c) 1999 - 2009 The PHP Group. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. The name "PHP" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact group@php.net.

4. Products derived from this software may not be called "PHP", nor may "PHP" appear in their name, without prior written permission from group@php.net. You may indicate that your software works in conjunction with PHP by saying "Foo for PHP" instead of calling it "PHP Foo" or "phpfoo"

5. The PHP Group may publish revised and/or new versions of the license from time to time. Each version will be given a distinguishing version number. Once covered code has been published under a particular version of the license, you may always continue to use it under the terms of that version. You may also choose to use such covered code under the terms of any subsequent version of the license published by the PHP Group. No one other than the PHP Group has the right to modify the terms applicable to covered code created under this License.

6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes PHP software, freely available from <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE PHP DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project, please see <http://www.php.net>.

PHP includes the Zend Engine, freely available at <http://www.zend.com>.

# Zlib

Copyright (C) 1995-2005 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1.The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.

2.Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.

3.This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org

Mark Adler madler@alumni.caltech.edu

# Index


**A**

access domains
  firewall 56, 61
  Panorama 392
accounts
  authentication profiles 62
  creating administrator 59
  username and password requirements 57
acknowledging alarms 85
active configuration, updating 37, 41
Active Directory
  configuring User-ID Agent 298
  installing User-ID Agent 297
  uninstalling and upgrading User-ID Agent 301
  User-ID Agent 296
active/active high availability 93
active/passive high availability 93
adding devices to Panorama 382
address groups, defining 230
addresses
  defining 227
  defining address groups 230
  defining group 230
  defining ranges 227
administrator
  accounts, about 56
  accounts, creating 59
  authentication options 56
  page lockout 61, 384, 391
  profiles, about 56
  roles, about 56
  roles, defining 58
agent
  configuring terminal server 301
  GlobalProtect 335
  setting up GlobalProtect 353
  User-ID 284
  using GlobalProtect 353
aggregate Ethernet interfaces
  configuring 141
  groups 141
aggregate groups 141
alarms
  acknowledged 85
  alarm icon 85
  log 68
  log settings 73
  making the icon visible 85
  thresholds 178
  unacknowledged 85
  viewing 85
allow list
  URL filtering profile 219
  wildcard patterns 218
anti-spyware profiles
  about 213
  defining 213
antivirus profiles
  defining 212
  settings 212
antivirus response pages 429
App-ID, requesting 277
Application Command Center (ACC), using 253
application exception policies 246
application exceptions 213
application groups, defining 238
application override policies
  about 205
  defining 205
applications
  ACC page 254
  categories 233, 437
  characteristics 234, 439
  custom with application override 205
  defining 233
  defining filters 238
  defining groups 238
  details 232
  exceptions 213
  filters 231
  identifying unknown 276
  response page 431
  searching 232
  sub category 233
  subcategories 437
  technologies 439
  updating threat definitions 55
applications list 254
App-Scope reports

change monitor report 258
network monitor report 261
summary report 257
threat map report 259, 260, 263
viewing 256
ARP entries
on L3 subinterfaces 144
on main L3 interfaces 138
on VLAN interfaces 134, 146
AS
BGP 154
description 154
audit configuration 49
authentication
GlobalProtect 336
IKE 312
LDAP 56
local database 56
options for administrator 56
RADIUS 56
remote 30, 49
sequence 67
Authentication Header (AH) 312
authentication profiles
about 62
Kerberos settings 67
LDAP settings 66
RADIUS settings 65
setting up 62
authentication sequences
about 68
setting up 68

**B**

backing up firewall configurations 417
BGP
about 153
configuring virtual router 163
description 154
redistribution profiles 158
virtual routers 160
block list
URL filtering profile 218
wildcard patterns 218
blocking, file profiles 220
botnet infected clients 267
botnets
about 267
reports 267
BrightCloud service 217
browsers, supported 27

**C**

candidate configuration
about 37, 41
saving and rolling back 37, 41
captive portal 62
comfort page 115, 434
configuring firewall for 290
defining policies 206
settings 207, 209
certificate authority (CA)
CRL 46
GlobalProtect 336
OCSP 46
trusted CA certificate 86
certificates
CRL 46
decryption 86
exporting 87
generating 87
GlobalProtect 336
importing 87
OCSP 46
Panorama server 87
renew 87
revoke 86
trusted CA 86
web 86
clear text traffic, and QoS 356
clients
botnet infected 267
downloading and activating
GlobalProtect 352
clock, setting 30, 49
committing
changes 25
options 25
Panorama 398
comparison of configurations 49
configuration audit 49
configuration bundle exports 417
configuration log
about 68
defining remote logging 72, 73, 74
viewing 267
configuration management 37, 41
configuration, sample VPN 320
content-id settings 43
conventions, typographical 15
CPU utilization 252
CRL 46
crossover cables 100
crypto profiles 318, 319
crypto profiles, about 312
custom group reports 272
custom reports 274
custom signatures
about 244
spyware 244
vulnerability 244

**D**

dashboard
firewall 252
data filtering

ACC page 255
data patterns 243
defining profiles 223
HIP matches on ACC page 255
list 255
pattern settings 224
profile settings 223, 225
profiles 223
profiles and patterns 224
viewing logs 69, 266
data patterns
adding new 240
data filtering profiles 224
defining 243
rules 240
data protection
adding 45
changing password 45
dead peer protection 315
decoders and actions 212
decryption policy 248
defining configuration templates 419
denial of service (Dos), profiles 225
deployment options
Layer 2 124
Layer 3 124
PPPoE 124
tap mode 125
virtual wire 120
deployment types 120
deployment, viewing information 416
deployments 120
device groups
adding 386
object assignment and sharing 396
selecting 396
using 396
device priority, HA 378
devices
adding 385
management 29
master 387
DHCP
firewall options 171
relay 171
servers 171
settings 171, 356
Diffie-Hellman (DH) group 312, 318
discard options, DOS profiles 180
disk utilization 252
DNS
servers 172
DNS proxy
about 173
settings 173
do not fragment (DF) 181
domain name 30
DoS
defining policies 209
profiles 225
protection profiles 225
downgrading software 53
duplex settings 128, 139, 149, 150
Duplicate Address Detection (DAD) 132, 136, 144
Dynamic Block Lists 242
dynamic updates
about 55
scheduling 56
dynamic URL timeout 44

**E**

editing settings on a page 24
email
scheduling report delivery 273
email notification settings
defining 83, 85
in logging profiles 247
Encapsulating Security Payload (ESP) 312
encrypting private keys and passwords 90
Ethernet interfaces, configuring 142
exchange mode 314
exports
certificates 86
configuration bundle 417
scheduling log 71

**F**

fail over 177
features and benefits 18
file blocking
defining profiles 220
profiles, defining 246
settings 220
file blocking page 431
filters
application 231, 238
sub category 231
FIPS 441
firewall
configuring WildFire settings 423
features and benefits 18
introduction 17
latitude and longitude 31
navigating the user interface 26
User-ID Agent 286
using the web interface 23
flood protection 175
flood, zone protection settings 178, 358
FTP server, saving logs to 71

**G**

gateway
about 335
GlobalProtect 335
setting up GlobalProtect 347
getting help 24

GlobalProtect
about 335
agents 335
authentication 336
connection process 336
downloading and activating clients 352
gateways 335
large scale VPN 323
portal backup 334
portals 335
response page 115
setting up 337
setting up agents 353
setting up gateways 347
setting up portals 341
using the agent 353
groups
aggregate interface 141
defining service 240
device 386

## H

HA interfaces 150
HA upgrading 51
HA1 and HA2 ports 101
hello interval, HA 378
help 24
high availability
about 93
active/active 93
active/passive 93
configuration notes 108
configuring 101
configuring on Panorama 377
defining interfaces 150
enabling 101
interfaces 150
notes about setting up 100
Panorama 377
rules for operation and failover 93
setting up 99
upgrading PAN-OS software 51
hold time 378
Host Information Profile (HIP)
HIP match log settings 73
match log 267
matches on ACC page 255
setting up 340
setting up objects 337
host name, defining 30, 49
HTML block pages 429

## I

ICMP flood 178, 179
identification, IKE 312
IKE
about 310, 312
AH 312
authentication 312
crypto profile settings 318
crypto profiles, about 312
dead peer protection 315
defining crypto profiles 318
DH group 312
ESP 312
exchange mode 314
identification 312
IKE gateways
about 312
setting up 175, 314
settings 314
installation, Panorama 372
interface management profiles 175
interfaces
aggregate groups 141
configuring aggregate Ethernet 141, 142
high availability 150
L2 subinterfaces 129, 149
L2, main 128
L3 subinterfaces 134
L3, main 130
summary of 127
tap 149
viewing status 128, 252
IPSec
AH 312
crypto profile settings 319
defining crypto profiles 319
DH group 312
ESP 312
large scale VPN overview 323
lifetime 312
number of tunnels 311
setting up tunnels 315
IPv6 176
IPv6 addresses 228

## K

Kerberos
administrator roles 56
configuring server settings 67
knowledge base 117

## L

L2 deployments, about 124
L2 interfaces
configuring 128
configuring subinterfaces 129
main 128
subinterfaces 129, 149
L3 deployments, about 124
L3 interfaces
configuring 130
configuring subinterfaces 134
loopback 146, 147
main 130

shared gateways 115
subinterfaces 134
Large Scale VPN
overview 323
latitude and longitude 31
Layer 130
LDAP
authentication 56
configuring server settings 66
licenses
installing 50, 86
open source 445
link groups, HA 106
link speed and duplex 128, 139, 149, 150
link state
setting 128, 139, 149, 150
viewing 252
local identification 314
lockout on Administrator's page 61, 384, 391
log destinations
about 70
email 83, 85
SNMP traps 75
syslog 76
log exports 71
log forwarding
defining profiles 247
profile settings 247
log page links 254
logs 266
about 68
alarms 68, 73
clearing 74
configuration 68
configuration settings 72
defining remote logging
for the configuration 72, 73, 74
for threat and traffic logs 246
HIP match 267
HIP match settings 73
links from ACC pages 254
managing 74
resolve hostname 265
saving to FTP server 71
scheduling exports 71
system 69
threat 69
viewing 264
viewing URL filtering 266
WildFire 70
loopback interfaces
defining 146, 147
management port 30
LSA 153

**M**

management interface
CLI 19
configuring 30, 49
options 19
Panorama 19
web 19
managing configurations 37, 41
master device 387
Master Key and Diagnostics page 90
MD5 162
memory utilization 252
MIBs 47, 75
modifying settings on a page 24
monitor profiles 177
multicast routing
about 154
settings 169
multiple virtual systems 31, 111, 113

**N**

NAT
defining policies 194
NAT64 196
policies 190
policy examples 195
types 192
navigation 26
Netflow
about 85
configuring 85
network profiles 174
network settings 30, 49
networking overview 120
next hop 157
NFS 373
external log storage 373
Panorama high availability 378
storage partitions 373
NIS servers 172
NSSA (not so stub area) 161
NT LAN Manager (NTLM) 208, 285
NTP servers 172

**O**

objects
overview 226
sharing in Panorama 395
OCSP 46
open source licenses 445
Open Virtual Machine Format (OVF) 372
OSPF
about 153
configuring virtual router 160
redistribution profiles 158
virtual routers 160

**P**

packet capture 266
accessing 266

APP-ID 278
capture files 278
configuring capture settings 278
profile setting 212, 214, 215, 216
taking captures 278
packet-based attack protection 175, 180
Panorama
ACC tab 382
access domains 392
accessing 382
adding devices 385
administrator account creation 389
administrator options 387
administrator roles 388
committing 398
configuration bundle exports 417
configuring IP address 32
dashboard 382
device tab 383
enabling access 32
hardware appliance 376
high availability 377
installing 372
interface description 382
logging 403
monitor tab 382
network tab 382
objects 395
objects tab 382
panorama tab 383
policies tab 382
server certificate 87
set up 371
shared policies, defining 382, 393
tab 383
tabs 382
templates 399
templates, configuring 401
upgrading software 417, 419
user account lockout 61, 384, 391
user interface 382
virtual appliance 372
PAN-OS software
upgrading 50, 61, 392
version 252
PAN-OS, upgrading software 50
passive hold time, HA 378
passive link state 106
passive/active high availability 93
password
data protection 45
encrypting 90
minimum password complexity 35
new 22
profiles 59
path groups, HA 379
PDF summary reports
creating 271, 273
designing 271
displaying 271
viewing 270
peer identification 314
Perfect Forward Security (PFS) 312
policies
about 183
about NAT 190
about policy based forwarding 199
about security 187
and virtual systems 111, 113
data patterns 243
defining address objects 227
defining captive portal 206
defining decryption 202
defining DoS 209
defining NAT 194
guidelines on defining 184
other policy objects 226
post rules 394
pre rules 393
QoS 359
shared 382, 393
specifying users and applications 186
types 183
virtual systems 110
policies shared global 393
policy based forwarding (PBF)
about 199
and monitor profiles 177
defining 199
portal
about 335
GlobalProtect 335
setting up GlobalProtect 341
post rules in policies 394
PPPoE
about 124
deployments 124
settings 131
pre rules in policies 393
private key, encrypting 90
profile groups, defining 246
profiles
about monitor 177
about security 211
anti-spyware 213
antivirus 212
antivirus, application exceptions 213
antivirus, decoders and actions 212
data filtering 223
defining log forwarding 247
file blocking 220, 246
flood protection 175
IKE crypto 318
IKE crypto profile settings 318
interface management 175
IPSec crypto 319
IPSec crypto profile settings 319
logging 246

network 174
packet-based attack 175
QoS 355, 358
reconnaissance detection 175
redistribution 158
security groups 211, 246
tunnel monitor 177
URL filtering 217
vulnerability protection 215, 217
zone protection 178, 358
proxy
DNS, about 173

## Q

QoS
classes 358, 359
clear text traffic 356
configuration instructions 355
egress settings 358
marking 190
policies 359
priority settings 358
profiles 355, 358
settings 190
settings for firewall interfaces 355
tunneled traffic 356
quality of service (QoS) 355

## R

RADIUS
authentication 56
authentication profiles 62
defining server settings 65
random early drop 178
rebooting the device 30, 39, 49
rebooting the firewall 39
reconnaissance detection 175
redistribution profiles
about 158
configuring 158
regions
about 230
policies 230
regular expressions, data patterns 240
rematching sessions 45
remote authentication 30, 49
rendezvous point 169
reports
App-Scope 256
creating custom group 272
custom 274
PDF summary 270
scheduling email delivery 273
top 50 273
user activity 272, 404
viewing 273
viewing WildFire 427
reports and logs
custom reports 274
identifying unknown applications 276
using the Application Command Center 253
using the dashboard 252
viewing App-Scope reports 256
viewing PDF summary reports 270
viewing reports 273, 274
Representational State Transfer (REST) 19
requesting support 117
required fields 26
resolve hostname 265
response pages
antivirus 115, 429
application block 115, 431
captive portal 115, 434
defining 115
file blocking 115, 431
file blocking continue 115
GlobalProtect portal help 115
GlobalProtect portal login 115
SSL certificate errors notify page 116
SSL certificate revoked notify 436
SSL decryption opt-out 116
types 89, 115
URL filtering continue and override 116, 434
response thresholds 178, 179
restarting the firewall 39
RIP
about 153
configuring virtual router 159
redistribution profiles 158
virtual routers 159
roles
about 56
defining administrator 58
rolling back a candidate configuration 37, 41
Router Advertisement 133, 137, 145
routing
about virtual routers and routing protocols 153
multicast 154
routing protocols
about 153
BGP 160
OSPF 160
RIP 159
virtual routers 153
rules
application exception policy 246
security policy 187

## S

sample VPN configuration 320
saving a candidate configuration 37, 41
schedules
configuration bundle exports 417
defining 246, 250
security

defining profile groups 211, 246
profile groups 246
security associations (SA) 311
security certificates 86
security policies
about 187
defining 187
security profile groups, defining 246
security profiles
about 211
actions 211
defining 246
security zones
about 151
defining 151
in NAT policies 194
in security policies 188
interfaces 151
sensitive information, protecting 45
servers
defining Kerberos 67
defining LDAP 66
defining RADIUS 65
defining syslog 76
service groups
defining 240
service groups, defining 240
services, defining 239
session browser 267
shared gateways
about 112
common interface 112
configuring 115
L3 interfaces 115
shared policy
defining 382, 393
master device 387
Shortest Path Tree (SPT) 170
signatures
custom 244
spyware 244
vulnerability 244
SNMP
community string 47
defining trap destinations 75
MIB setup 47
MIBs 75
SNMP trap destinations
defining 75
in logging profiles 247
software
downgrading 53
upgrading 50, 61, 392, 417, 419
upgrading Panorama 419
version 252
source-specific multicast (SSM) 171
speed, link 128, 139, 149, 150
SSL
decryption policies 246
decryption rule settings 200, 203
defining decryption policies 202
tech notes reference 202
SSL VPNs
about 355
authentication profiles 62
comfort page 116
local user database 64
split tunnels 349
storage partitions 373
Panorama 373
sub category
application 233
filtering 231
support information 117
support information, viewing 117
supported browsers 27
SYN flood 178
syslog servers
custom syslog fields 77
defining 76
in logging profiles 247
system log
about 69
viewing 267
system settings 115


## T


tables, using in web interface 26
tags
on L2 subinterfaces 129, 140
on L3 subinterfaces 135
on virtual wires 126
tap interfaces, defining 149
tap mode
deployment option 125
description 125
terminal services agent (TS Agent) 301
threat list 255
threat log 248
about 69
defining remote logging 246
viewing 266
threats
ACC list 255
updating definitions 55
thresholds, alarm 178
time
setting 30, 49
zone 31
traffic log 247
defining remote logging 246
viewing 266
Transport Layer Security (TLS) 66
TS Agent
configuring firewall to support 301
configuring on the terminal server 303
installing 302

uninstalling 307
upgrading 302
tunnel interfaces 315
tunnel monitor
fail over 177
profiles 177
wait-recover 177
tunneled traffic, and QoS 356
tunnels
about VPN 311
number of IPSec 311
setting up 315
split for SSL VPNs 349
types of deployments 120
typographical conventions 15


## U


UDP flood 179
unknown applications
identifying 276
requesting App-ID 277
taking action 277
unnumbered loopback interfaces 146
upgrading
Panorama software 417, 419
PAN-OS software 50, 61, 392
schedules 56
threat and application definitions 55
with high availability 51
URL filtering
ACC page 254
continue and override response page 116, 434
defining profiles 217
dynamic categorization 217
list 254
override settings 44
profile settings 217
response pages 116
viewing log 266
viewing logs 69
user account lockout 61, 391
user database, SSL VPN 64
user interface navigation 26
User-ID Agent
captive portal configuration 290
configuring firewall 286
configuring for Active Directory 298
for Active Directory, about 296
installing for Active Directory 297
overview 281
uninstalling and upgrading for Active Directory 301
username and password requirements 57


## V


version, software 252
viewing
devices 387
logs 264
session browser 267
session information 267
virtual firewall 363
virtual routers
configuring 155, 157, 169
defining 155
multicast settings 169
next hop 157
routing protocols 153
runtime statistics 171
virtual systems
about 110
and security zones 111
communications among 111
defining 110, 113, 115
defining multiple 113
enabling 31
enabling multiple 31
internal traffic flow 111
multiple 111
policies 110
security zones 110
shared gateway common interface 112
shared gateways 112
virtual wire 120
defining 125
interfaces 141
interfaces, configuring 138
VLANs
interfaces, defining 143
L2 interfaces 143
VM-Series firewall
installation 363
installing 366
limitations 364
overview 363
requirements 364
troubleshooting 370
VMware ESX(i) 372
VPN
about 310
IPSec and IKE crypto profiles 312
large scale VPN and Dynamic Routing Protocols 333
large scale VPN deployment 324
large scale VPN overview 323
sample configuration 320
setting up tunnels 311
SSL, about 355
VPN tunnels
about 311
IKE 311
manual security keys 311
securing 311
setting up 313, 315
vSphere 372
vulnerability protection profiles 215, 217

## W

web interface
committing changes 25
navigation 26
required fields 26
supported browsers 27
using 23
using tables 26
wildcard
custom URL categories 242
patterns for allow and block lists 218
WildFire
configure forwarding 424
configuring firewall settings 423
configuring settings on the portal 426
dashboard 425
log descriptions 425
setup tasks 423
subscription 422
supported file types 422
using the portal 425
viewing reports 427
WINS servers 172

## X

XML API 19

## Z

zones
defining 151
in NAT policies 194
in security policies 188
protection profiles 178, 358

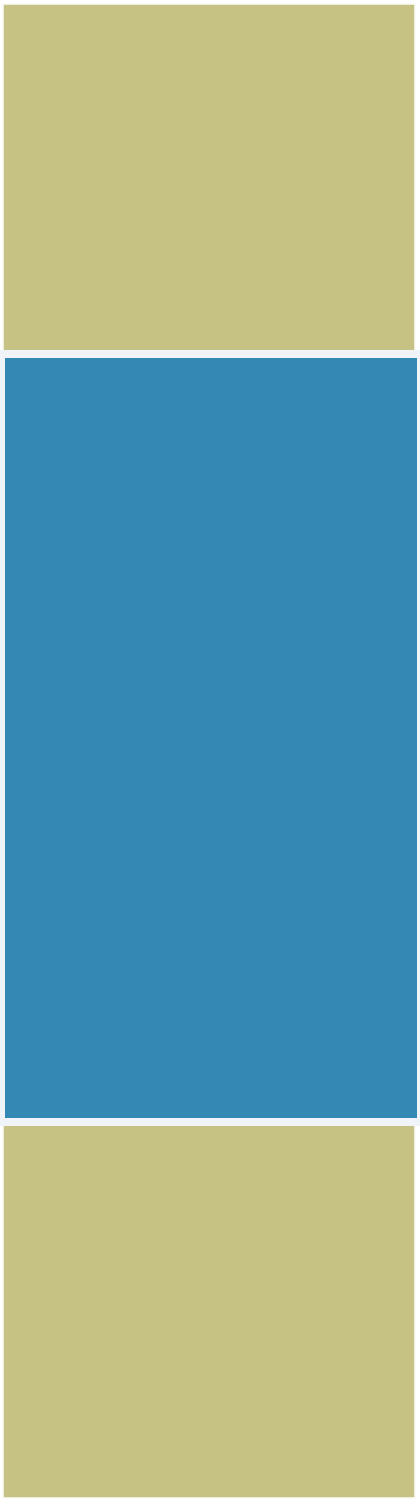

Palo Alto Networks

# PAN-OS® Administrator's Guide
Version 6.0

## Contact Information

**Corporate Headquarters:**
Palo Alto Networks
4401 Great America Parkway
Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-us

## About this Guide

This guide provides the concepts and solutions to help you get the most out of your Palo Alto Networks next-generation firewalls. For additional information, refer to the following resources:

- For start-to-finish instruction on how to set up a new firewall, refer to the Palo Alto Networks Getting Started Guide.
- For information on the additional capabilities and for instructions on configuring the features on the firewall, refer to https://www.paloaltonetworks.com/documentation.
- For access to the knowledge base, discussion forums, and videos, refer to https://live.paloaltonetworks.com.
- For contacting support, for information on support programs, to manage your account or devices, or to open a support case, refer to https://www.paloaltonetworks.com/support/tabs/overview.html.
- For the most current PAN-OS 6.0 release notes, go to https://www.paloaltonetworks.com/documentation/60/pan-os/pan-os-6-0-release-notes.html.

To provide feedback on the documentation, please write to us at: documentation@paloaltonetworks.com.

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2007–2016 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at http://www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

**Revision Date: March 10, 2016**

# Table of Contents


## Getting Started . . . . . 11

Integrate the Firewall into Your Management Network . . . . . 12
- Set Up Management Access to the Firewall . . . . . 13
- Activate Firewall Services . . . . . 20

Create the Security Perimeter . . . . . 26
- Security Perimeter Overview . . . . . 27
- Basic Interface Deployments . . . . . 28
- About Network Address Translation (NAT) . . . . . 30
- About Security Policies . . . . . 31
- Set Up Interfaces and Zones . . . . . 36
- Configure NAT Policies . . . . . 39
- Set Up Basic Security Policies . . . . . 43

Enable Basic Threat Prevention Features . . . . . 49
- Enable WildFire . . . . . 50
- Scan Traffic for Threats . . . . . 52
- Control Access to Web Content . . . . . 57

Best Practices for Completing the Firewall Deployment . . . . . 60

## Device Management . . . . . 61

Management Interfaces . . . . . 62
- Use the Web Interface . . . . . 63
- Use the Command Line Interface (CLI) . . . . . 68
- Use the XML API . . . . . 70

Manage Firewall Administrators . . . . . 72
- Administrative Roles . . . . . 72
- Administrative Authentication . . . . . 73
- Create an Administrative Account . . . . . 73

Reference: Web Interface Administrator Access . . . . . 77
- Web Interface Access Privileges . . . . . 77
- Panorama Web Interface Access . . . . . 100

## Certificate Management . . . . . 103

Keys and Certificates . . . . . 104

Certificate Revocation . . . . . 106
- Certificate Revocation List (CRL) . . . . . 106
- Open Certificate Status Protocol (OCSP) . . . . . 106

Certificate Deployment . . . . . 108

Set Up Verification for Certificate Revocation Status . . . . . 109
- Configure an OCSP Responder . . . . . 109
- Configure Revocation Status Verification of Certificates Used for User/Device Authentication . . . 110
- Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption . . . . . 111

Configure the Master Key . . . . . 113


© Palo Alto Networks, Inc.

Obtain Certificates . . . 114
Create a Self-Signed Root CA Certificate . . . 114
Generate a Certificate on the Firewall . . . 115
Import a Certificate and Private Key . . . 116
Obtain a Certificate from an External CA . . . 117
Configure a Certificate Profile . . . 119
Revoke and Renew Certificates . . . 122
Revoke a Certificate . . . 122
Renew a Certificate . . . 122
Secure Keys with a Hardware Security Module . . . 123
Set up Connectivity with an HSM . . . 124
Encrypt a Master Key Using an HSM . . . 130
Store Private Keys on an HSM . . . 132
Manage the HSM Deployment . . . 134

High Availability . . . 135
HA Overview . . . 136
HA Concepts . . . 137
HA Modes . . . 137
HA Links and Backup Links . . . 137
Device Priority and Preemption . . . 139
Failover Triggers . . . 140
HA Timers . . . 141
Set Up Active/Passive HA . . . 143
Prerequisites for Active/Passive HA . . . 144
Configuration Guidelines for Active/Passive HA . . . 145
Configure Active/Passive HA . . . 147
Define HA Failover Conditions . . . 153
Verify Failover . . . 154
HA Resources . . . 155

Reports and Logging . . . 157
Use the Dashboard . . . 158
Use the Application Command Center . . . 160
ACC Risk Level . . . 161
ACC Charts . . . 162
ACC Detail Pages . . . 164
Use the ACC . . . 165
Use App-Scope . . . 166
Summary Report . . . 167
Change Monitor Report . . . 168
Threat Monitor Report . . . 170
Threat Map Report . . . 171
Network Monitor Report . . . 172
Traffic Map Report . . . 173
Take Packet Captures . . . 174


© Palo Alto Networks, Inc.

Monitor the Firewall . . . 176
Monitor Applications and Threats . . . 177
Monitor Log Data . . . 178
Monitor the Dashboard . . . 182
View Reports . . . 183
Forward Logs to External Services . . . 184
Define Remote Logging Destinations . . . 185
Enable Log Forwarding . . . 192
Monitor the Firewall Using SNMP . . . 194
Monitor the Firewall Using NetFlow . . . 196
Identify Firewall Interfaces in External Monitoring Systems . . . 197
Manage Reporting . . . 199
About Reports . . . 200
View Reports . . . 201
Disable Predefined Reports . . . 202
Generate Custom Reports . . . 203
Generate Botnet Reports . . . 209
Manage PDF Summary Reports . . . 211
Generate User/Group Activity Reports . . . 213
Manage Report Groups . . . 214
Schedule Reports for Email Delivery . . . 215
Syslog Field Descriptions . . . 216
Traffic Logs . . . 216
Threat Logs . . . 218
HIP Match Logs . . . 222
Config Logs . . . 223
System Logs . . . 224
Syslog Severity . . . 225
Custom Log/Event Format . . . 225
Escape Sequences . . . 225

User-ID . . . 227

User-ID Overview . . . 228
User-ID Concepts . . . 230
Group Mapping . . . 230
User Mapping . . . 230
Enable User-ID . . . 234
Map Users to Groups . . . 235
Map IP Addresses to Users . . . 237
Configure User Mapping Using the Windows User-ID Agent . . . 238
Configure User Mapping Using the PAN-OS Integrated User-ID Agent . . . 244
Configure User-ID to Receive User Mappings from a Syslog Sender . . . 247
Map IP Addresses to User Names Using Captive Portal . . . 257
Configure User Mapping for Terminal Server Users . . . 263
Send User Mappings to User-ID Using the XML API . . . 271
Configure a Firewall to Share User Mapping Data with Other Firewalls . . . 272
Enable User- and Group-Based Policy . . . 274
Verify the User-ID Configuration . . . 277

© Palo Alto Networks, Inc.

App-ID . . . . . 281
App-ID Overview . . . . . 282
Manage Custom or Unknown Applications . . . . . 283
Best Practices for Using App-ID in Policy . . . . . 285
Applications with Implicit Support . . . . . 286
About Application Level Gateways . . . . . 289
Disable the SIP Application-level Gateway (ALG) . . . . . 290

Threat Prevention . . . . . 291
License the Threat Prevention Features . . . . . 292
About Threat Prevention Licenses . . . . . 292
Obtain and Install Licenses . . . . . 292
About Security Profiles . . . . . 294
Set Up Security Profiles and Policies . . . . . 299
Set Up Antivirus, Anti-spyware, and Vulnerability Protection . . . . . 299
Set Up Data Filtering . . . . . 302
Set Up File Blocking . . . . . 306
Prevent Brute Force Attacks . . . . . 308
Customize the Action and Trigger Conditions for a Brute Force Signature . . . . . 309
Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions . . . . . 311
Passive DNS Collection . . . . . 314
Content Delivery Network Infrastructure for Dynamic Updates . . . . . 315
Threat Prevention Resources . . . . . 317

Decryption . . . . . 319
Decryption Overview . . . . . 320
Decryption Concepts . . . . . 321
Keys and Certificates for Decryption Policies . . . . . 322
SSL Forward Proxy . . . . . 324
SSL Inbound Inspection . . . . . 326
SSH Proxy . . . . . 327
Decryption Exceptions . . . . . 328
Decryption Port Mirroring . . . . . 329
Configure SSL Forward Proxy . . . . . 330
Configure SSL Inbound Inspection . . . . . 334
Configure SSH Proxy . . . . . 336
Configure Decryption Exceptions . . . . . 337
Exclude Traffic From Decryption . . . . . 337
Exclude a Server From Decryption . . . . . 338
Configure Decryption Port Mirroring . . . . . 339


© Palo Alto Networks, Inc.

## URL Filtering . . . . . 341

URL Filtering Overview . . . . . 342
- URL Filtering Vendors . . . . . 342
- Interaction Between App-ID and URL Categories . . . . . 342

URL Filtering Concepts . . . . . 344
- URL Categories . . . . . 344
- URL Filtering Profiles . . . . . 345
- URL Category as Policy Match Criteria . . . . . 349

PAN-DB Categorization Workflow . . . . . 351
- PAN-DB URL Categorization Components . . . . . 351
- PAN-DB URL Categorization Workflow . . . . . 352

Configure URL Filtering . . . . . 354
- Enable URL Filtering . . . . . 355
- Determine URL Filtering Policy Requirements . . . . . 357
- Define Website Controls . . . . . 365

Enable Safe Search Enforcement . . . . . 368
- Block Search Results that are not Using Strict Safe Search Settings . . . . . 368
- Enable Transparent Safe Search Enforcement . . . . . 372

URL Filtering Use Case Examples . . . . . 376
- Use Case: Control Web Access . . . . . 377
- Use Case: Use URL Categories in Policy . . . . . 381

Troubleshoot URL Filtering . . . . . 383
- Problems Activating PAN-DB . . . . . 383
- PAN-DB Cloud Connectivity Issues . . . . . 384
- URLs Classified as Not-Resolved . . . . . 385
- Incorrect Categorization . . . . . 386
- URL Database Out of Date . . . . . 387

## Quality of Service . . . . . 389

QoS Overview . . . . . 390

QoS Concepts . . . . . 392
- QoS for Applications and Users . . . . . 392
- QoS Profile . . . . . 393
- QoS Classes . . . . . 394
- QoS Policy . . . . . 395
- QoS Egress Interface . . . . . 396
- QoS Cleartext and Tunneled Traffic . . . . . 397

Configure QoS . . . . . 398

Configure QoS for a Virtual System . . . . . 403

QoS Use Case Examples . . . . . 410
- QoS for a Single User . . . . . 411
- QoS for Voice and Video Applications . . . . . 414


© Palo Alto Networks, Inc.

VPNs . . . . . 417
VPN Deployments . . . . . 418
Site-to-Site VPN Overview . . . . . 419
Site-to-Site VPN Concepts . . . . . 420
IKE Gateway . . . . . 420
Tunnel Interface . . . . . 420
Tunnel Monitoring . . . . . 421
Internet Key Exchange (IKE) for VPN . . . . . 421
Set Up Site-to-Site VPN . . . . . 425
Set up an IKE Gateway . . . . . 426
Define Cryptographic Profiles . . . . . 428
Set up an IPSec Tunnel . . . . . 430
Set up Tunnel Monitoring . . . . . 434
Test VPN Connectivity . . . . . 436
Interpret VPN Error Messages . . . . . 437
Site-to-Site VPN Quick Configs . . . . . 439
Site-to-Site VPN with Static Routing . . . . . 440
Site-to-Site VPN with OSPF . . . . . 445
Site-to-Site VPN with Static and Dynamic Routing . . . . . 451
Large Scale VPN (LSVPN) . . . . . 459
LSVPN Overview . . . . . 460
Create Interfaces and Zones for the LSVPN . . . . . 461
Enable SSL Between GlobalProtect LSVPN Components . . . . . 463
About Certificate Deployment . . . . . 463
Deploy Server Certificates to the GlobalProtect LSVPN Components . . . . . 463
Configure the Portal to Authenticate Satellites . . . . . 467
Configure GlobalProtect Gateways for LSVPN . . . . . 469
Prerequisite Tasks . . . . . 469
Configure the Gateway . . . . . 469
Configure the GlobalProtect Portal for LSVPN . . . . . 473
Prerequisite Tasks . . . . . 473
Configure the Portal . . . . . 473
Define the Satellite Configurations . . . . . 474
Prepare the Satellite Device to Join the LSVPN . . . . . 478
Verify the LSVPN Configuration . . . . . 481
LSVPN Quick Configs . . . . . 482
Basic LSVPN Configuration with Static Routing . . . . . 483
Advanced LSVPN Configuration with Dynamic Routing . . . . . 486
Networking . . . . . 489
Interface Deployments . . . . . 490
Virtual Wire Deployments . . . . . 490
Layer 2 Deployments . . . . . 493
Layer 3 Deployments . . . . . 493
Tap Mode Deployments . . . . . 494


© Palo Alto Networks, Inc.

Configure a Virtual Router . . . . . 495
Configure Static Routes . . . . . 497
Configure RIP . . . . . 499
Configure OSPF . . . . . 501
OSPF Concepts . . . . . 501
Configure OSPF . . . . . 503
Configure OSPFv3 . . . . . 508
Configure OSPF Graceful Restart . . . . . 511
Confirm OSPF Operation . . . . . 512
Configure BGP . . . . . 515

© Palo Alto Networks, Inc.

# Getting Started

The following sections provide detailed steps to help you deploy a new Palo Alto Networks next-generation firewall. They provide details for integrating a new firewall into your network and configuring basic security policies and threat prevention features.

After you perform the basic configuration steps required to integrate the firewall into your network, you can use the rest of the topics in this guide to help you deploy the comprehensive enterprise security platform features as necessary to address your network security needs.

- Integrate the Firewall into Your Management Network
- Create the Security Perimeter
- Enable Basic Threat Prevention Features
- Best Practices for Completing the Firewall Deployment

# Integrate the Firewall into Your Management Network

The following topics describe how to perform the initial configuration steps that are necessary to integrate a new firewall into the management network and deploy it in a basic security configuration.

The following topics describe how to integrate a single Palo Alto Networks next-generation firewall into your network. For details on how to deploy a pair of firewalls in a high availability configuration, read the information in HA documentation before proceeding.

- Set Up Management Access to the Firewall
- Activate Firewall Services


© Palo Alto Networks, Inc.

# Set Up Management Access to the Firewall

All Palo Alto Networks firewalls provide an out-of-band management port (MGT) that you can use to perform the firewall administration functions. By using the MGT port, you separate the management functions of the firewall from the data processing functions, safeguarding access to the firewall and enhancing performance. When using the web interface, you must perform all initial configuration tasks from the MGT port even if you plan to use an in-band port for managing your device going forward.

Some management tasks, such as retrieving licenses and updating the threat and application signatures on the firewall require access to the Internet. If you do not want to enable external access to your MGT port, you will need to either set up a data port to provide access to required external services or plan to manually upload updates regularly.

The following sections provide instructions for setting up management access to the firewall:

- Determine Your Management Strategy
- Perform Initial Configuration
- Set Up Network Access for External Services

## Determine Your Management Strategy

The Palo Alto Networks firewall can be configured and managed locally or it can be managed centrally using Panorama, the Palo Alto Networks centralized security management system. If you have six or more firewalls deployed in your network, use Panorama to achieve the following benefits:

- Reduce the complexity and administrative overhead in managing configuration, policies, software and dynamic content updates. Using device groups and templates on Panorama, you can effectively manage device specific configuration locally on a device and enforce shared policies across all devices or device groups.
- Aggregate data from all managed firewalls and gain visibility across all the traffic on your network. The Application Command Center (**ACC**) on Panorama provides a single glass pane for unified reporting across all the firewalls, allowing you to centrally analyze, investigate and report on network traffic, security incidents and administrative modifications.

The procedures in this document describe how to manage the firewall using the local web interface. If you want to use Panorama for centralized management, after you complete the instructions in the Perform Initial Configuration section of this guide and verify that the firewall can establish a connection to Panorama. From that point on you can use Panorama to configure your firewall centrally.

## Perform Initial Configuration

By default, the firewall has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other firewall configuration tasks. You must perform these initial configuration tasks either from the MGT interface, even if you do not plan to use this interface for your firewall management, or using a direct serial connection to the console port on the device.

| Set Up Network Access to the Firewall | | |
|---|---|---|
| Step 1 | Gather the required information from your network administrator. | • IP address for MGT port<br>• Netmask<br>• Default gateway<br>• DNS server address |
| Step 2 | Connect your computer to the firewall. | You can connect to the firewall in one of the following ways:<br>• Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1). Wait a few minutes for the boot-up sequence to complete; when the device is ready, the prompt changes to the name of the firewall, for example `PA-500 login`.<br>• Connect an RJ-45 Ethernet cable from your computer to the MGT port on the firewall. From a browser, go to `https://192.168.1.1`. Note that you may need to change the IP address on your computer to an address in the 192.168.1.0 network, such as 192.168.1.2, in order to access this URL. |
| Step 3 | When prompted, log in to the firewall. | You must log in using the default username and password (admin/admin). The firewall will begin to initialize. |
| Step 4 | Configure the MGT interface. | 1. Select **Device > Setup > Management** and then click the Edit icon in the Management Interface Settings section of the screen. Enter the **IP Address**, **Netmask**, and **Default Gateway**.<br>2. Set the **Speed** to **auto-negotiate**.<br>3. Select which management services to allow on the interface.<br>Best Practice:<br>Make sure **Telnet** and **HTTP** are not selected because these services use plaintext and are not as secure as the other services.<br>4. Click **OK**. |
| Step 5 | (Optional) Configure general firewall settings. | 1. Select **Device > Setup > Management** and click the Edit icon in the General Settings section of the screen.<br>2. Enter a **Hostname** for the firewall and enter your network **Domain** name. The domain name is just a label; it will not be used to join the domain.<br>3. Enter the **Latitude** and **Longitude** to enable accurate placement of the firewall on the world map.<br>4. Click **OK**. |

 © Palo Alto Networks, Inc.

| Set Up Network Access to the Firewall (Continued) | | |
|---|---|---|
| Step 6 | Configure DNS, time and date settings.<br><br>You must manually configure at least one DNS server on the firewall or it will not be able to resolve hostnames; it will not use DNS server settings from another source, such as an ISP. | 1. Select **Device > Setup > Services** and click the Edit icon in the Services section of the screen.<br>2. On the **Services** tab, enter the IP address of your **Primary DNS Server** and optionally your **Secondary DNS Server**.<br>3. To use the virtual cluster of time servers on the Internet, enter the hostname **pool.ntp.org** as the **Primary NTP Server** or add the IP address of your **Primary NTP Server** and optionally your **Secondary NTP Server**.<br>4. Click **OK** to save your settings. |
| Step 7 | Set a secure password for the admin account. | 1. Select **Device > Administrators.**<br>2. Select the **admin** role.<br>3. Enter the current default password and the new password.<br>4. Click **OK** to save your settings. |
| Step 8 | Commit your changes.<br><br>When the configuration changes are saved, you will lose connectivity to the web interface because the IP address will have changed. | Click **Commit**. The device may take up to 90 seconds to save your changes.<br> |
| Step 9 | Connect the firewall to your network. | 1. Disconnect the firewall from your computer.<br>2. Connect the MGT port to a switch port on your management network using an RJ-45 Ethernet cable. Make sure that the switch port you cable the firewall to is configured for auto-negotiation. |
| Step 10 | Open an SSH management session to the firewall. | Using a terminal emulation software, such as PuTTY, launch an SSH session to the firewall using the new IP address you assigned to it. |
| Step 11 | Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server, in one of the following ways:<br>• If you do not want to allow external network access to the MGT interface, you will need to set up a data port to retrieve required service updates. Continue to Set Up Network Access for External Services.<br>• If you do plan to allow external network access to the MGT interface, verify that you have connectivity and then proceed to Activate Firewall Services. | If you cabled your MGT port for external network access, verify that you have access to and from the firewall by using the ping utility from the CLI. Make sure you have connectivity to the default gateway, DNS server, and the Palo Alto Networks Update Server as shown in the following example:<br>`admin@PA-200> ping host updates.paloaltonetworks.com`<br>`PING updates.paloaltonetworks.com (67.192.236.252) 56(84) bytes of data.`<br>`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=40.5 ms`<br>`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=53.6 ms`<br>`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=79.5 ms`<br>After you have verified connectivity, press Ctrl+C to stop the pings. |

## Set Up Network Access for External Services

By default, the firewall uses the MGT interface to access remote services, such as DNS servers, content updates, and license retrieval. If you do not want to enable external network access to your management network, you must set up a data port to provide access to these required external services.

This task requires familiarity with firewall interfaces, zones, and policies. For more information on these topics, see Create the Security Perimeter.

| Set Up a Data Port for Access to External Services | | |
|---|---|---|
| Step 1 | Decide which port you want to use for access to external services and connect it to your switch or router port. | The interface you use will need to have a static IP address. |
| Step 2 | Log in to the web interface. | Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page. |
| Step 3 | (Optional) The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and zones). If you do not plan to use this virtual wire configuration, you must manually delete the configuration to prevent it from interfering with other interface settings you define. | You must delete the configuration in the following order:<br>1. To delete the default security policy, select **Policies > Security**, select the rule, and click **Delete**.<br>2. Next, delete the default virtual wire by selecting **Network > Virtual Wires**, selecting the virtual wire and clicking **Delete**.<br>3. To delete the default trust and untrust zones, select **Network > Zones**, select each zone and click **Delete**.<br>4. Finally, delete the interface configurations by selecting **Network > Interfaces** and then select each interface (ethernet1/1 and ethernet1/2) and click **Delete**.<br>5. **Commit** the changes. |

 © Palo Alto Networks, Inc.

| Set Up a Data Port for Access to External Services (Continued) | |
|---|---|
| Step 4 Configure the interface. | 1. Select **Network > Interfaces** and select the interface that corresponds to the port you cabled in Step 1.<br>2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**.<br>4. In the Zone dialog, define a **Name** for new zone, for example L3-trust, and then click **OK**.<br>5. Select the **IPv4** tab, select the **Static** radio button, and click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.254/24.<br>6. Select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**.<br>7. Enter a **Name** for the profile, such as allow_ping, and then select the services you want to allow on the interface. These services provide management access to the device, so only select the services that correspond to the management activities you want to allow on this interface. For example, if you plan to use the MGT interface for device configuration tasks through the web interface or CLI, you would not want to enable HTTP, HTTPS, SSH, or Telnet so that you could prevent unauthorized access through this interface. For the purposes of allowing access to the external services you probably only need to enable **Ping** and then click **OK**.<br>8. To save the interface configuration, click **OK**. |

**Set Up a Data Port for Access to External Services (Continued)**

Step 5 Because the firewall uses the MGT interface by default to access the external services it requires, you must change the interface the firewall uses to send these requests by editing the service routes.

1. Select **Device > Setup > Services > Service Route Configuration**.

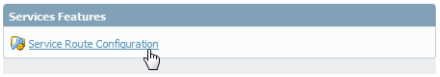


For the purposes of activating your licenses and getting the most recent content and software updates, you will want to change the service route for **DNS**, **Palo Alto Updates**, **URL Updates**, and **WildFire**.

2. Click the **Customize** radio button, and select one of the following:
   - For a predefined service, select **IPv4** or **IPv6** and click the link for the service for which you want to modify the **Source Interface** and select the interface you just configured.

     If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.
   - To create a service route for a custom destination, select **Destination**, and click **Add**. Enter a **Destination** name and select a **Source Interface**. If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.
3. Click **OK** to save the settings.
4. Repeat steps 2-3 above for each service route you want to modify.
5. **Commit** your changes.

 © Palo Alto Networks, Inc.

**Set Up a Data Port for Access to External Services (Continued)**

Step 6 Configure an external-facing interface and an associated zone and then create security and NAT policy rules to allow the firewall to send service requests from the internal zone to the external zone:

1. Select **Network > Interfaces** and then select your external-facing interface. Select **Layer3** as the **Interface Type**, **Add** the **IP** address (on the **IPv4** or **IPv6** tab), and create the associated **Security Zone** (on the **Config** tab), such as l3-untrust. You do not need to set up management services on this interface.
2. To set up a security rule that allows traffic from your internal network to the Palo Alto Networks update server and external DNS servers, select **Policies > Security** and click **Add**. For the purposes of initial configuration, you can create a simple rule that allows all traffic from l3-trust to l3-untrust as follows:

3. If you are using a private IP address on the internal-facing interface, you will need to create a source NAT rule to translate the address to a publicly routable address. Select **Policies > NAT** and then click **Add**. At a minimum you must define a name for the rule (**General** tab), specify a source and destination zone, l3-trust to l3-untrust in this case (**Original Packet** tab), and define the source address translation settings (**Translated Packet** tab) and then click **OK**. For more information on NAT, see Configure NAT Policies.

4. **Commit** your changes.

| | | |
|---|---|---|
| Step 7 | Verify that you have connectivity from the data port to the external services, including the default gateway, DNS server, and the Palo Alto Networks Update Server.<br><br>After you verify you have the required network connectivity, continue to Activate Firewall Services. | Launch the CLI and use the ping utility to verify that you have connectivity. Keep in mind that by default pings are sent from the MGT interface, so in this case you must specify the source interface for the ping requests as follows:<br>`admin@PA-200> ping source 192.168.1.254 host updates.paloaltonetworks.com`<br>`PING updates.paloaltonetworks.com (67.192.236.252) from 192.168.1.254 : 56(84) bytes of data.`<br>`64 bytes from 67.192.236.252: icmp_seq=1 ttl=242 time=56.7 ms`<br>`64 bytes from 67.192.236.252: icmp_seq=2 ttl=242 time=47.7 ms`<br>`64 bytes from 67.192.236.252: icmp_seq=3 ttl=242 time=47.6 ms`<br>`^C`<br><br>After you have verified connectivity, press Ctrl+C to stop the pings. |

# Activate Firewall Services

Before you can begin using the firewall to secure your network, you must register it and activate the licenses for the services you have purchased. In addition, you should ensure that you are running the appropriate version of PAN-OS as described in the following sections:

- Register With Palo Alto Networks
- Activate Licenses
- Manage Content Updates
- Install Software Updates

## Register With Palo Alto Networks

| Register the Firewall | | |
|---|---|---|
| Step 1 | Log in to the web interface. | Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page. |
| Step 2 | Locate your serial number and copy it to the clipboard. | On the **Dashboard**, locate your **Serial Number** in the General Information section of the screen. |
| Step 3 | Go to the Palo Alto Networks Support site. | In a new browser tab or window, go to https://support.paloaltonetworks.com. |
| Step 4 | Register the device. The way you register depends on whether you already have a login to the support site. | • If this is the first Palo Alto Networks device you are registering and you do not yet have a login, click **Register** on the right side of the page. To register, you must provide your sales order number or customer ID, and the serial number of your firewall (which you can paste from your clipboard) or the authorization code you received with your order. You will also be prompted to set up a username and password for access to the Palo Alto Networks support community.<br>• If you already have a support account, log in and then click **My Devices**. Scroll down to Register Device section at the bottom of the screen and enter the serial number of your firewall (which you can paste from your clipboard), your city and postal code and then click **Register Device**. |

## Activate Licenses

Before you can start using your firewall to secure the traffic on your network, you must activate the licenses for each of the services you purchased. Available licenses and subscriptions include the following:

- Threat Prevention—Provides antivirus, anti-spyware, and vulnerability protection.

 © Palo Alto Networks, Inc.

- Decryption Port Mirroring—Provides the ability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as NetWitness or Solera—for archiving and analysis.
- URL Filtering—In order to create policy rules based on dynamic URL categories, you must purchase and install a subscription for one of the supported URL filtering databases: PAN-DB or BrightCloud. For more information about URL filtering, see Control Access to Web Content.
- **Virtual Systems**—This license is required to enable support for multiple virtual systems on PA-2000 and PA-3000 Series firewalls. In addition, you must purchase a Virtual Systems license if you want to increase the number of virtual systems beyond the base number provided by default on PA-4000 Series, PA-5000 Series, and PA-7050 firewalls (the base number varies by platform). The PA-500, PA-200, and VM-Series firewalls do not support virtual systems.
- WildFire—Although basic WildFire support is included as part of the Threat Prevention license, the WildFire subscription service provides enhanced services for organizations that require immediate coverage for threats, enabling sub-hourly WildFire signature updates, advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet), as well as the ability to upload files using the WildFire API. A WildFire subscription is also required if your firewalls will be forwarding files to a private WF-500 WildFire appliance.
- GlobalProtect—Provides mobility solutions and/or large-scale VPN capabilities. By default, you can deploy a single GlobalProtect portal and gateway (without HIP checks) without a license. However, if you want to deploy multiple gateways, you must purchase a portal license (one-time, permanent license). If you want to use host checks you will also need gateway licenses (subscription) for each gateway.

| **Activate Licenses** | | |
|---|---|---|
| Step 1 | Locate the activation codes for the licenses you purchased. | When you purchased your subscriptions you should have received an email from Palo Alto Networks customer service listing the activation code associated with each subscription. If you cannot locate this email, contact customer support to obtain your activation codes before you proceed. |
| Step 2 | Launch the web interface and go to the license page. | Select **Device > Licenses**. |
| Step 3 | Activate each license you purchased.<br> If your firewall does not have Internet access from the management port, you can manually download your license files from the support site and upload them to your firewall using the **Manually upload license key** option. | 1. Select **Activate feature using authorization code**.<br>2. When prompted, enter the **Authorization Code** and then click **OK**.<br>3. Verify that the license was successfully activated. For example, after activating the WildFire license, you should see that the license is valid:<br> |

## Manage Content Updates

In order to stay ahead of the changing threat and application landscape, Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to the Palo Alto Networks devices. The devices access the web resources in the CDN to perform various App-ID and Content-ID functions. By default, the devices use the management port to access the CDN infrastructure for application updates, threat and antivirus signature updates, BrightCloud and PAN-DB database updates and lookups, and access to the Palo Alto Networks WildFire Cloud. To ensure that you are always protected from the latest threats (including those that have not yet been discovered), you must ensure that you keep your devices up-to-date with the latest updates published by Palo Alto Networks.

The following content updates are available, depending on which subscriptions you have:

Although you can manually download and install content updates at any time, as a best practice you should schedule updates to occur automatically.

- **Antivirus**—Includes new and updated antivirus signatures, including signatures discovered by the WildFire cloud service. You must have a Threat Prevention subscription to get these updates. New antivirus signatures are published daily.
- **Applications**—Includes new and updated application signatures. This update does not require any additional subscriptions, but it does require a valid maintenance/support contract. New application updates are published weekly.
- **Applications and Threats**—Includes new and updated application and threat signatures. This update is available if you have a Threat Prevention subscription (and you get it instead of the Applications update). New Applications and Threats updates are published weekly.
- **GlobalProtect Data File**—Contains the vendor-specific information for defining and evaluating host information profile (HIP) data returned by GlobalProtect agents. You must have a GlobalProtect portal and GlobalProtect gateway license in order to receive these updates. In addition, you must create a schedule for these updates before GlobalProtect will function.
- **BrightCloud URL Filtering**—Provides updates to the BrightCloud URL Filtering database only. You must have a BrightCloud subscription to get these updates. New BrightCloud URL database updates are published daily. If you have a PAN-DB license, scheduled updates are not required as devices remain in-sync with the servers automatically.
- **WildFire**—Provides near real-time malware and antivirus signatures created as a result of the analysis done by the WildFire cloud service. Without the subscription, you must wait 24 to 48 hours for the signatures to roll into the Applications and Threat update.

If your firewall does not have Internet access from the management port, you can download content updates from the Palo Alto Networks Support Site (https://support.paloaltonetworks.com) and then Upload them to your firewall.

If your firewall is deployed behind existing firewalls or proxy servers, access to these external resources might be restricted using access control lists that allow the firewall to only access a hostname or an IP address. In such cases, to allow access to the CDN, set the update server address to use the hostname staticupdates.paloaltonetworks.com or the IP address `199.167.52.15`

 © Palo Alto Networks, Inc.

| Download the Latest Databases | | |
|---|---|---|
| Step 1 | Verify that the firewall points to the CDN infrastructure. | Select **Device > Setup > Services**.<br>• As a best practice, set the **Update Server** to access `updates.paloaltonetworks.com`. This allows the firewall to receive content updates from the server to which it is closest in the CDN infrastructure.<br>• (Optional) If the firewall has restricted access to the internet, set the update server address to use the hostname staticupdates.paloaltonetworks.com or the IP address 199.167.52.15. For additional security, select Verify Update Server Identity. The firewall verifies that the server from which the software or content package is download has an SSL certificate signed by a trusted authority. |
| Step 2 | Launch the web interface and go to the Dynamic Updates page. | Select **Device > Dynamic Updates**. |
| Step 3 | Check for the latest updates. | Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available:<br>• **Download**—Indicates that a new update file is available. Click the link to begin downloading the file directly to the firewall. After successful download, the link in the **Action** column changes from **Download** to **Install**.<br> <br>You cannot download the antivirus database until you have installed the Application and Threats database.<br>• **Upgrade**—Indicates that there is a new version of the BrightCloud database available. Click the link to begin the download and installation of the database. The database upgrade begins in the background; when completed a check mark displays in the **Currently Installed** column. Note that if you are using PAN-DB as your URL filtering database you will not see an upgrade link because the PAN-DB database automatically stays in sync with the server.<br> <br>To check the status of an action, click **Tasks** (on the lower right-hand corner of the window).<br>• **Revert**—Indicates that the corresponding software version has been downloaded previously. You can choose to revert to the previously installed version of the update. |

**Download the Latest Databases (Continued)**

Step 4 Install the updates.

Installation can take up to 20 minutes on a PA-200, PA-500, or PA-2000 device and up to two minutes on a PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7050, or VM-Series firewall.

Click the **Install** link in the **Action** column. When the installation completes, a check mark displays in the **Currently Installed** column.

Step 5 Schedule each update.

Repeat this step for each update you want to schedule.

Best Practice:

Stagger the update schedules because the firewall can only download one update at a time. If you schedule the updates to download during the same time interval, only the first download will succeed.

1. Set the schedule of each update type by clicking the **None** link.

2. Specify how often you want the updates to occur by selecting a value from the **Recurrence** drop-down. The available values vary by content type (WildFire updates are available **Every 15 minutes**, **Every 30 minutes** or **Every Hour** whereas all other content types can be scheduled for **Daily** or **Weekly** update).
3. Specify the **Time** and (or, minutes past the hour in the case of WildFire), if applicable depending on the **Recurrence** value you selected, **Day** of the week that you want the updates to occur.
4. Specify whether you want the system to **Download And Install** the update (best practice) or **Download Only**.
5. In rare instances, errors in content updates may be found. For this reason, you may want to delay installing new updates until they have been released for a certain number of hours. You can specify how long after a release to wait before performing a content update by entering the number of hours to wait in the **Threshold (Hours)** field.
6. Click **OK** to save the schedule settings.
7. Click **Commit** to save the settings to the running configuration.

## Install Software Updates

When installing a new firewall, it is a good idea to upgrade to the latest software update (or to the update version recommended by your reseller or Palo Alto Networks Systems Engineer) to take advantage of the latest fixes and security enhancements. Note that before updating the software, you should first make sure you have the latest content updates as detailed in the previous section (the release notes for a software update specify the minimum content update versions that are supported in the release).

 © Palo Alto Networks, Inc.

| Update PAN-OS | | |
|---|---|---|
| Step 1 | Launch the web interface and go to the Software page. | Select **Device > Software**. |
| Step 2 | Check for software updates. | Click **Check Now** to check for the latest updates. If the value in the **Action** column is **Download** it indicates that an update is available. |
| Step 3 | Download the update.<br><br>If your firewall does not have Internet access from the management port, you can download the software update from the Palo Alto Networks Support Site (https://support.paloaltonetworks.com). You can then manually **Upload** them to your firewall. | Locate the version you want and then click **Download**. When the download completes, the value in the Action column changes to **Install**.<br> |
| Step 4 | Install the update. | 1. Click **Install**.<br>2. Reboot the firewall:<br>• If you are prompted to reboot, click **Yes**.<br> <br>• If you are not prompted to reboot, select **Device > Setup > Operations** and click **Reboot Device** in the Device Operations section of the screen.<br> |

# Create the Security Perimeter

The following topics provide basic steps for configuring the firewall interfaces, defining zones, and setting up a basic security policy:

- Security Perimeter Overview
- Configure NAT Policies
- Set Up Basic Security Policies


© Palo Alto Networks, Inc.

## Security Perimeter Overview

Traffic must pass through the firewall in order for the firewall to manage and control it. Physically, traffic enters and exits the firewall through *interfaces*. The firewall decides how to act on a packet based on whether the packet matches a *security policy*. At the most basic level, the security policy must identify where the traffic came from and where it is going. On a Palo Alto Networks next-generation firewall, security policies are applied between zones. A *zone* is a grouping of interfaces (physical or virtual) that provides an abstraction for an area of trust for simplified policy enforcement. For example, in the following topology diagram, there are three zones: Trust, Untrust, and DMZ. Traffic can flow freely within a zone, but traffic will not be able to flow between zones until you define a security policy that allows it.

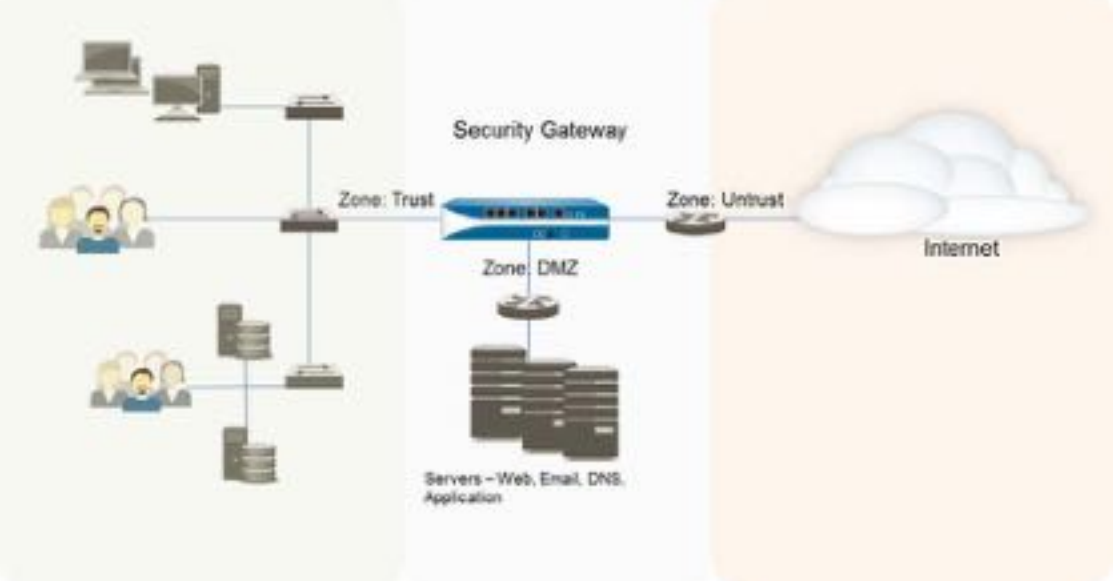


The following sections describe the components of the security perimeter and provide steps for configuring the firewall interfaces, defining zones, and setting up a basic security policy that allows traffic from your internal zone to the Internet and to the DMZ. By initially creating a basic policy like this, you will be able to analyze the traffic running through your network and use this information to define more granular policies for safely enabling applications while preventing threats.

- Basic Interface Deployments
- About Network Address Translation (NAT)
- About Security Policies

# Basic Interface Deployments

All Palo Alto Networks next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, enabling you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports. The following sections provide basic information on each type of deployment.

- Virtual Wire Deployments
- Layer 2 Deployments
- Layer 3 Deployments

For more detailed deployment information, refer to Designing Networks with Palo Alto Networks Firewalls.

## Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together. By using a virtual wire, you can install the firewall in any network environment without reconfiguring adjacent devices. If necessary, a virtual wire can block or allow traffic based on the virtual LAN (VLAN) tag values. You can also create multiple subinterfaces and classify traffic according to an IP Address (address, range, or subnet), VLAN, or a combination of the two.

By default, the virtual wire (named *default-vwire*) binds Ethernet ports 1 and 2 and allows all untagged traffic. Choose this deployment to simplify installation and configuration and/or avoid configuration changes to surrounding network devices.

A virtual wire is the default configuration, and should be used only when no switching or routing is needed. If you do not plan to use the default virtual wire, you should manually delete the configuration before proceeding with interface configuration to prevent it from interfering with other interface settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up a Data Port for Access to External Services.

## Layer 2 Deployments

In a Layer 2 deployment, the firewall provides switching between two or more interfaces. Each group of interfaces must be assigned to a VLAN object in order for the firewall to switch between them. The firewall will perform VLAN tag switching when Layer 2 subinterfaces are attached to a common VLAN object. Choose this option when switching is required.

For more information on Layer 2 deployments, refer to the Layer 2 Networking Tech Note and/or the Securing Inter VLAN Traffic Tech Note.

 © Palo Alto Networks, Inc.

## Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between ports. An IP address must be assigned to each interface and a virtual router must be defined to route the traffic. Choose this option when routing is required.

You must assign an IP address to each physical Layer 3 interface you configure. You can also create logical subinterfaces for each physical Layer 3 interface that allows you to segregate the traffic on the interface based on VLAN tag (when VLAN trunking is in use) or by IP address, for example for multi-tenancy.

In addition, because the firewall must route traffic in a Layer 3 deployment, you must configure a virtual router. You can configure the virtual router to participate with dynamic routing protocols (BGP, OSPF, or RIP) as well as adding static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that are not shared between virtual routers, enabling you to configure different routing behaviors for different interfaces.

The configuration example in this chapter illustrates how to integrate the firewall into your Layer 3 network using static routes. For information on other types of routing integrations, refer to the following documents:

- How to Configure OSPF Tech Note
- How to Configure BGP Tech Note

## About Network Address Translation (NAT)

When you use private IP addresses within your internal networks, you must use network address translation (NAT) in order to translate the private addresses to public addresses that can be routed on external networks. In PAN-OS, you create NAT policy rules that instruct the firewall which packets need translation and how to do the translation. The firewall supports both source address and/or port translation and destination address and/or port translation. For more details about the different types of NAT rules, refer to the Understanding and Configuring NAT Tech Note.

It is important to understand the way the firewall applies the NAT and security policies in order to determine what policies you need based on the zones you have defined. Upon ingress, the firewall inspects a packet to see if it matches any of the NAT rules that have been defined, based on source and/or destination zone. It then evaluates and applies any security rules that match the packet based on the original (pre-NAT) source and destination addresses. Finally, it translates the source and/or destination port numbers for any matching NAT rules upon egress. This distinction is important, because it means that the firewall determines what zone a packet is destined for based on the address on the packet, not on the placement of the device based on its internally assigned address.

 © Palo Alto Networks, Inc.

# About Security Policies

Security policies protect network assets from threats and disruptions and aid in optimally allocating network resources for enhancing productivity and efficiency in business processes. On the Palo Alto Networks firewall, security policies determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service. By default, intra-zone traffic (that is traffic within the same zone, for example *from* trust *to* trust), is allowed. Traffic between different zones (or inter-zone traffic) is blocked until you create a security policy to allow the traffic.

Security policies are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log, if logging is enabled for that rule. The logging options are configurable for each rule, and can for example be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.


- Components of a Security Policy
- Policy Best Practices
- About Policy Objects
- About Security Profiles


## Components of a Security Policy

The security policy construct permits a combination of the required and optional components listed below.

| | Field | Description |
|---|---|---|
| **Required Fields** | **Name** | A label that supports up to 31 characters, used to identify the rule. |
| | **Source Zone** | The zone from which the traffic originates. |
| | **Destination Zone** | The zone at which the traffic terminates. If you use NAT, make sure to always reference the post-NAT zone. |
| | **Application** | The application which you wish to control. The firewall uses App-ID, the traffic classification technology, to identify traffic on your network. App-ID provides application control and visibility in creating security policies that block unknown applications, while enabling, inspecting, and shaping those that are allowed. |
| | **Action** | Specifies an *Allow* or *Deny* action for the traffic based on the criteria you define in the rule. |
| **Optional Fields** | **Tag** | A keyword or phrase that allows you to filter security rules. This is handy when you have defined many rules and wish to then review those that are tagged with a particular keyword, for example *Inbound to DMZ*. |
| | **Description** | A text field, up to 255 characters, used to describe the rule. |

| | Field | Description (Continued) |
|---|---|---|
| | **Source IP Address** | Define host IP or FQDN, subnet, named groups, or country-based enforcement. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address). |
| | **Destination IP Address** | The location or destination for the traffic. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address). |
| | **User** | The user or group of users for whom the policy applies. You must have User-ID enabled on the zone. To enable User-ID, see User-ID Overview. |
| | **URL Category** | Using the URL Category as match criteria allows you to customize security profiles (antivirus, anti-spyware, vulnerability, file-blocking, Data Filtering, and DoS) on a per-URL-category basis. For example, you can prevent.exe file download/upload for URL categories that represent higher risk while allowing them for other categories. This functionality also allows you to attach schedules to specific URL categories (allow social-media websites during lunch & after-hours), mark certain URL categories with QoS (financial, medical, and business), and select different log forwarding profiles on a per-URL-category-basis.<br>Although you can manually configure URL categories on your device, to take advantage of the dynamic URL categorization updates available on the Palo Alto Networks firewalls, you must purchase a URL filtering license.<br> To block or allow traffic based on URL category, you must apply a URL Filtering profile to the security policy rules. Define the URL Category as *Any* and attach a URL Filtering profile to the security policy. See Create Security Rules for information on using the default profiles in your security policy and see Control Access to Web Content for more details. |
| | **Service** | Allows you to select a Layer 4 (TCP or UDP) port for the application. You can choose *any*, specify a port, or use *application-default* to permit use of the standards-based port for the application. For example, for applications with well- known port numbers such as DNS, the *application-default* option will match against DNS traffic only on TCP port 53. You can also add a custom application and define the ports that the application can use.<br> For inbound allow rules (for example, from untrust to trust), using *application-default* prevents applications from running on unusual ports and protocols. Application-default is the default option; while the device still checks for all applications on all ports, with this configuration, applications are only allowed on their standard ports/protocols. |
| | **Security Profiles** | Provide additional protection from threats, vulnerabilities, and data leaks. Security profiles are only evaluated for rules that have an *allow* action. |
| | **HIP Profile** (for GlobalProtect) | Allows you to identify clients with Host Information Profile (HIP) and then enforce access privileges. |

 © Palo Alto Networks, Inc.

| | Field | Description (Continued) |
|---|---|---|
| | **Options** | Allow you to define logging for the session, log forwarding settings, change Quality of Service (QoS) markings for packets that match the rule, and schedule when (day and time) the security rule should be in effect. |

## Policy Best Practices

The task of safely enabling Internet access and preventing misuse of web access privileges, and exposure to vulnerabilities and attacks is a continuous process. The key principle when defining policy on the Palo Alto Networks firewall is to use a positive enforcement approach. Positive enforcement implies that you selectively allow what is required for day-to-day business operations as opposed to a negative enforcement approach where you would selectively block everything that is not allowed. Consider the following suggestions when creating policy:

- If you have two or more zones with identical security requirements, combine them into one security rule.
- The ordering of rules is crucial to ensure the best match criteria. Because policy is evaluated top down, the more specific policy must precede the ones that are more general, so that the more specific rule is not *shadowed*. The term shadow refers to a rule that is not evaluated or is skipped because it is placed lower in the policy list. When the rule is placed lower, it is not evaluated because the match criteria was met by another rule that preceded it, thereby shadowing the rule from policy evaluation.
- To restrict and control access to inbound applications, in the security policy, explicitly define the port that the service/application will be listening on.
- Logging for broad allow rules—for example access to well known servers like DNS—can generate a lot of traffic. Hence it is not recommended unless absolutely necessary.
- By default, the firewall creates a log entry at the end of a session. However, you can modify this default behavior and configure the firewall to log at the start of the session. Because this significantly increases the log volume, logging at session start is recommended only when you are troubleshooting an issue. Another alternative for troubleshooting without enabling logging at session start is to use the session browser (**Monitor > Session Browser**) to view the sessions in real time.

## About Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With policy objects that are a collective unit, you can reference the object in security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

You can create the following policy objects on the firewall:

| Policy Object | Description |
|---|---|
| Address/Address Group, Region | Allow you to group specific source or destination addresses that require the same policy enforcement. The address object can include an IPv4 or IPv6 address (single IP, range, subnet) or the FQDN. Alternatively, a region can be defined by the latitude and longitude coordinates or you can select a country and define an IP address or IP range. You can then group a collection of address objects to create an *address group* object.<br>You can also use dynamic address groups to dynamically update IP addresses in environments where host IP addresses change frequently. |
| User/User Group | Allow you to create a list of users from the local database or an external database and group them. |
| Application Group and Application Filter | An *Application Filter* allows you to filter applications dynamically. It allows you to filter, and save a group of applications using the attributes defined in the application database on the firewall. For example, you can filter by one or more attributes—category, sub-category, technology, risk, characteristics—and save your application filter. With an application filter, when a PAN-OS content update occurs, any new applications that match your filter criteria are automatically added to your saved application filter.<br>An *Application Group* allows you to create a static group of specific applications that you wish to group together for a group of users or for a particular service. |
| Service/Service Groups | Allows you to specify the source and destination ports and protocol that a service can use. The firewall includes two pre-defined services—service-http and service-https—that use TCP ports 80 and 8080 for HTTP, and TCP port 443 for HTTPS. You can however, create any custom service on any TCP/UDP port of your choice to restrict application usage to specific ports on your network (in other words, you can define the default port for the application).<br> To view the standard ports used by an application, in **Objects > Applications** search for the application and click the link. A succinct description displays. |

Some examples of *address* and *application* policy objects are shown in the security policies that are included in Create Security Rules. For information on the other policy objects, see Enable Basic Threat Prevention Features.

## About Security Profiles

While security policies enable you to allow or deny traffic on your network, security profiles help you define an *allow but scan* rule, which scan allowed applications for threats. When traffic matches the allow rule defined in the security policy, the security profile(s) that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.

Security profiles are not used in the match criteria of a traffic flow. The security profile is applied to scan traffic after the application or category is allowed by the security policy.

The different types of security profiles that can be attached to security policies are: Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking, and Data Filtering. The firewall provides default security profiles that you can use out of the box to begin protecting your network from threats. See Create Security Rules for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, you can create custom profiles. See Scan Traffic for Threats for more

 © Palo Alto Networks, Inc.

information.

# Set Up Interfaces and Zones

The following sections provide information on configuring interfaces and zones:

- Plan the Deployment
- Configure Interfaces and Zones

## Plan the Deployment

Before you begin configuring interfaces and zones, take some time to plan the zones you will need based on the different usage requirements within your organization. In addition, you should gather all of the configuration information you will need ahead of time. At a basic level, you should plan which interfaces will belong to which zones. For Layer 3 deployments you'll also need to obtain the required IP addresses and network configuration information from your network administrator, including information on how to configure the routing protocol or static routes required for the virtual router configuration. The example in this chapter will be based on the following topology:

**Figure: Layer 3 Topology Example**

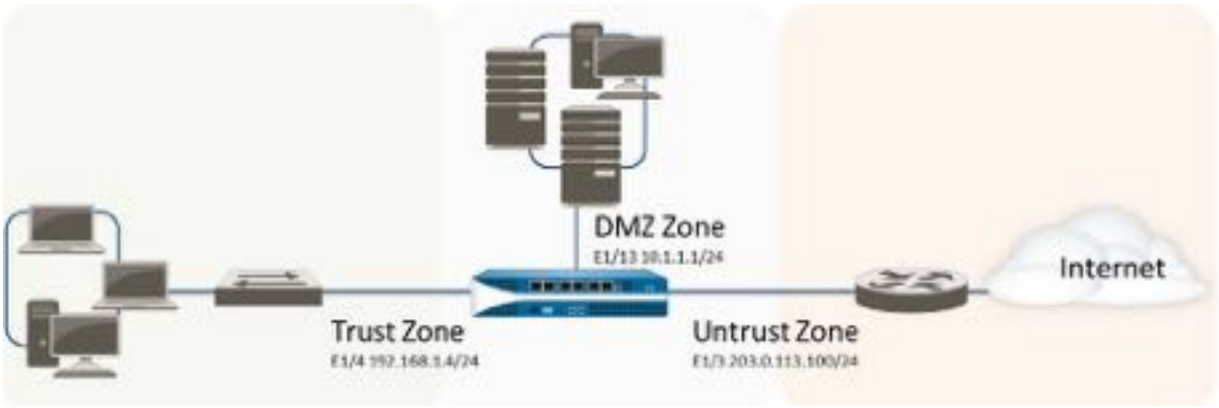


The following table shows the information we will use to configure the Layer 3 interfaces and their corresponding zones as shown in the sample topology.

| Zone | Deployment Type | Interface(s) | Configuration Settings |
|---|---|---|---|
| Untrust | L3 | Ethernet1/3 | **IP address:** 203.0.113.100/24<br>**Virtual router:** default<br>**Default route:** 0.0.0.0/0<br>**Next hop:** 203.0.113.1 |
| Trust | L3 | Ethernet1/4 | **IP address:** 192.168.1.4/24<br>**Virtual router:** default |
| DMZ | L3 | Ethernet1/13 | **IP address:** 10.1.1.1/24<br>**Virtual router:** default |

 © Palo Alto Networks, Inc.

## Configure Interfaces and Zones

After you plan your zones and the corresponding interfaces, you can configure them on the device. The way you configure each interface depends on your network topology.

The following procedure shows how to configure a Layer 3 deployment as depicted in Figure: Layer 3 Topology Example.

The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and virtual router). If you do not plan to use the default virtual wire, you must manually delete the configuration and commit the change before proceeding to prevent it from interfering with other settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up a Data Port for Access to External Services.

| Set Up Interfaces and Zones | | |
|---|---|---|
| Step 1 | Configure a default route to your Internet router. | 1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog.<br>2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and enter the route in the **Destination** field (for example, 0.0.0.0/0).<br>3. Select the **IP Address** radio button in the **Next Hop** field and then enter the IP address and netmask for your Internet gateway (for example, 203.0.113.1).<br>4. Click **OK** twice to save the virtual router configuration. |
| Step 2 | Configure the external interface (the interface that connects to the Internet). | 1. Select **Network > Interfaces** and then select the interface you want to configure. In this example, we are configuring Ethernet1/3 as the external interface.<br>2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.<br>3. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Untrust, and then click **OK**.<br>4. In the **Virtual Router** drop-down, select **default**.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 208.80.56.100/24.<br>6. To enable you to ping the interface, select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. Enter a **Name** for the profile, select **Ping** and then click **OK**.<br>7. To save the interface configuration, click **OK**. |

| Set Up Interfaces and Zones (Continued) | | |
|---|---|---|
| Step 3 | Configure the interface that connects to your internal network.<br><br>In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT. See Configure NAT Policies. | 1. Select **Network > Interfaces** and select the interface you want to configure. In this example, we are configuring Ethernet1/4 as the internal interface.<br>2. Select **Layer3** from the **Interface Type** drop down.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Trust, and then click **OK**.<br>4. Select the same Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 4 | Configure the interface that connects to the DMZ. | 1. Select the interface you want to configure.<br>2. Select **Layer3** from the **Interface Type** drop down. In this example, we are configuring Ethernet1/13 as the DMZ interface.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example DMZ, and then click **OK**.<br>4. Select the Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 5 | Save the interface configuration. | Click **Commit**. |
| Step 6 | Cable the firewall. | Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment. |
| Step 7 | Verify that the interfaces are active.<br><br> 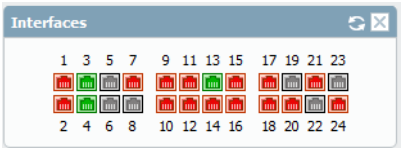 | From the web interface, select **Network > Interfaces** and verify that icon in the Link State column is green. You can also monitor link state from the **Interfaces** widget on the **Dashboard**. |

 © Palo Alto Networks, Inc.

## Configure NAT Policies

Based on the example topology we used to create the interfaces and zones, there are three NAT policies we need to create as follows:

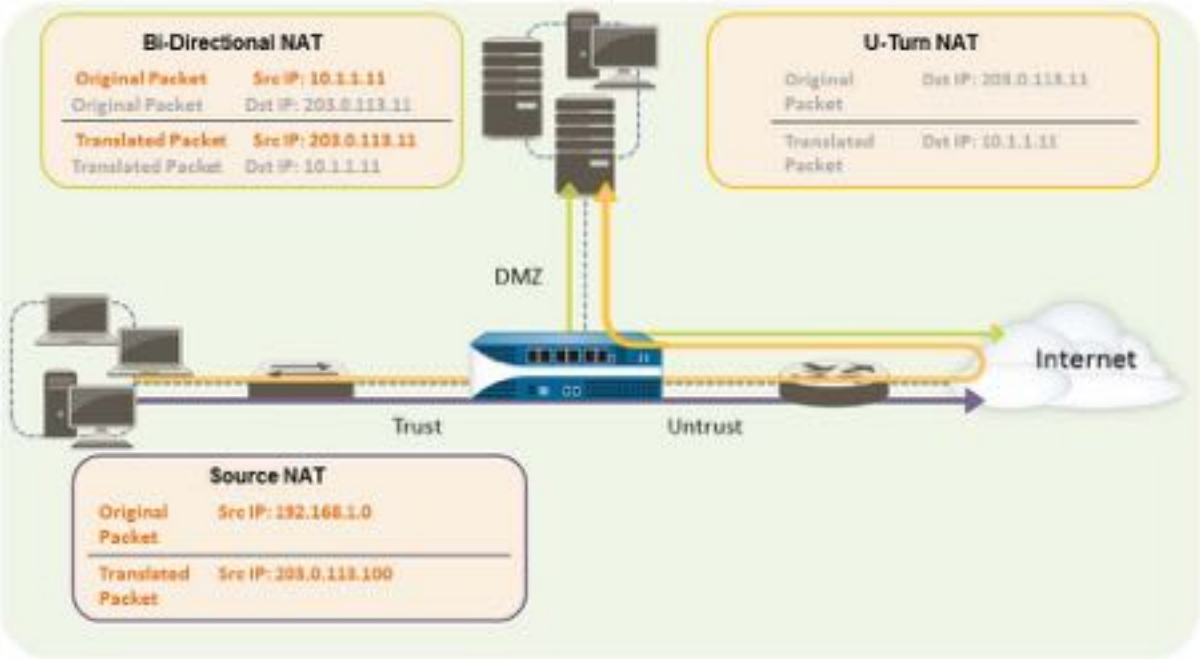


- To enable the clients on the internal network to access resources on the Internet, the internal 192.168.1.0 addresses will need to be translated to publicly routable addresses. In this case, we will configure source NAT, using the egress interface address, 203.0.113.100, as the source address in all packets that leave the firewall from the internal zone. See Translate Internal Client IP Addresses to your Public IP Address for instructions.
- To enable clients on the internal network to access the public web server in the DMZ zone, we will need to configure a NAT rule that redirects the packet from the external network, where the original routing table lookup will determine it should go based on the destination address of 203.0.113.11 within the packet, to the actual address of the web server on the DMZ network of 10.1.1.11. To do this you must create a NAT rule from the trust zone (where the source address in the packet is) to the untrust zone (where the original destination address is) to translate the destination address to an address in the DMZ zone. This type of destination NAT is called *U-Turn NAT*. See Enable Clients on the Internal Network to Access your Public Servers for instructions.
- To enable the web server—which has both a private IP address on the DMZ network and a public-facing address for access by external users—to both send and receive requests, the firewall must translate the incoming packets from the public IP address to the private IP address and the outgoing packets from the private IP address to the public IP address. On the firewall, you can accomplish this with a single bi-directional static source NAT policy. See Enable Bi-Directional Address Translation for your Public-Facing Servers.

## Translate Internal Client IP Addresses to your Public IP Address

When a client on your internal network sends a request, the source address in the packet contains the IP address for the client on your internal network. If you use private IP address ranges internally, the packets from the client will not be able to be routed on the Internet unless you translate the source IP address in the packets leaving the network into a publicly routable address. On the firewall you can do this by configuring a source NAT policy that translates the source address and optionally the port into a public address. One way to do this is to translate the source address for all packets to the egress interface on your firewall as shown in the following procedure.

**Configure Source NAT**

**Step 1** Create an address object for the external IP address you plan to use.

1. From the web interface, select **Objects > Addresses** and then click **Add**.
2. Enter a **Name** and optionally a **Description** for the object.
3. Select **IP Netmask** from the **Type** drop down and then enter the IP address of the external interface on the firewall, 203.0.113.100 in this example.
4. To save the address object, click **OK**.

Best Practice:

Although you do not have to use address objects in your policies, it is a best practice because it simplifies administration by allowing you to make updates in one place rather than having to update every policy where the address is referenced.

**Step 2** Create the NAT policy.

1. Select **Policies > NAT** and click **Add**.
2. Enter a descriptive **Name** for the policy.
3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop down.

4. On the **Translated Packet** tab, select **Dynamic IP And Port** from the **Translation Type** drop-down in the **Source Address Translation** section of the screen and then click **Add**. Select the address object you just created.

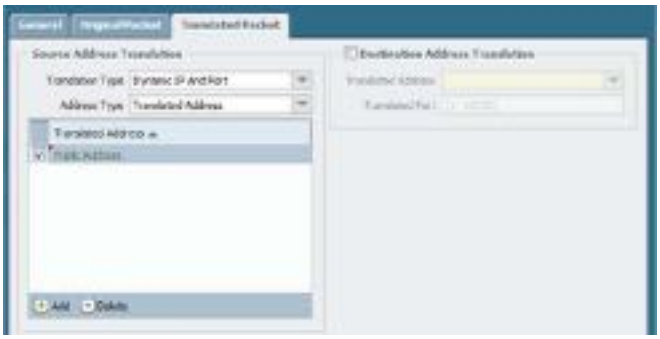

5. Click **OK** to save the NAT policy.

**Step 3** Save the configuration. Click **Commit**.

 © Palo Alto Networks, Inc.

## Enable Clients on the Internal Network to Access your Public Servers

When a user on the internal network sends a request for access to the corporate web server in the DMZ, the DNS server will resolve to the public IP address. When processing the request, the firewall will use the original destination in the packet (the public IP address) and route the packet to the egress interface for the untrust zone. In order for the firewall to know that it must translate the public IP address of the web server to an address on the DMZ network when it receives requests from users on the trust zone, you must create a destination NAT rule that will enable the firewall to send the request to the egress interface for the DMZ zone as follows.

**Configure U-Turn NAT**

Step 1 Create an address object for the web server.

1. From the web interface, select **Objects > Addresses** and then click **Add**.
2. Enter a **Name** and optionally a **Description** for the object.
3. Select **IP Netmask** from the **Type** drop down and then enter the public IP address of the web server, 203.0.113.11 in this example.
4. To save the address object, click **OK**.

Step 2 Create the NAT policy.

1. Select **Policies > NAT** and click **Add**.
2. Enter a descriptive **Name** for the NAT rule.
3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop down.
4. In the **Destination Address** section, click **Add** and select the address object you created for your public web server.

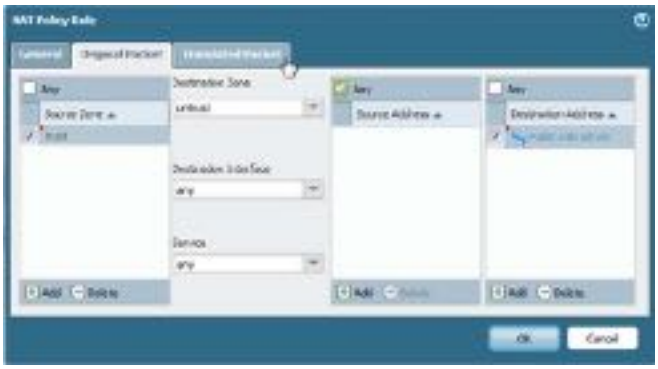

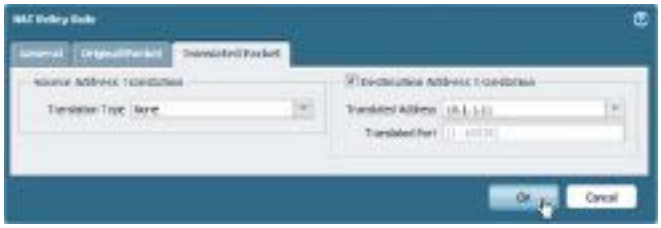

5. On the **Translated Packet** tab, select the **Destination Address Translation** check box and then enter the IP address that is assigned to the web server interface on the DMZ network, 10.1.1.11 in this example.
6. Click **OK** to save the NAT policy.

Step 3 Save the configuration. Click **Commit**.

## Enable Bi-Directional Address Translation for your Public-Facing Servers

When your public-facing servers have private IP addresses assigned on the network segment where they are physically located, you will need a source NAT rule for translating the source address of the server to the external address upon egress. You do this by creating a static NAT rule that instructs the firewall to translate the internal source address, 10.1.1.11, to the external web server address, 203.0.113.11 in our example. However, in the case of a public-facing server, the server must both be able to send packets and receive them. In this case, you need a reciprocal policy that will translate the public address that will be the destination IP address in incoming packets from users on the Internet into the private address to enable the firewall to properly route the packet to your DMZ network. On the firewall you do this by creating a bi-directional static NAT policy as described in the following procedure.

**Configure Bi-Directional NAT**

**Step 1** Create an address object for the web server's internal IP address.

1. From the web interface, select **Objects > Addresses** and then click **Add**.
2. Enter a **Name** and optionally a **Description** for the object.
3. Select **IP Netmask** from the **Type** drop down and then enter the IP address of the web server on the DMZ network, 10.1.1.11 in this example.
4. To save the address object, click **OK**.

If you did not already create an address object for the public address of your web server you should also create that object now.

**Step 2** Create the NAT policy.

1. Select **Policies > NAT** and click **Add**.
2. Enter a descriptive **Name** for the NAT rule.
3. On the **Original Packet** tab, select the zone you created for your DMZ in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop down.
4. In the **Source Address** section, click **Add** and select the address object you created for your internal web server address.

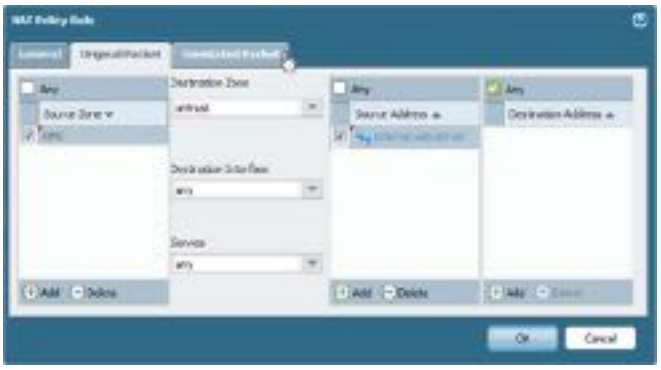

5. On the **Translated Packet** tab, select **Static IP** from the **Translation Type** drop down in the **Source Address Translation** section and then select the address object you created for your external web server address from the **Translated Address** drop down.
6. In the **Bi-directional** field, select **Yes**.
7. Click **OK** to save the NAT policy.

**Step 3** Save the configuration. Click **Commit**.

 © Palo Alto Networks, Inc.

# Set Up Basic Security Policies

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Captive Portal, Denial of Service, and Zone protection policies. All these different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. This section covers basic security policies and the default security profiles:

- Create Security Rules
- Test Your Security Policies
- Monitor the Traffic on Your Network

## Create Security Rules

Security policies reference security zones and enable you to allow, restrict, and track traffic on your network. Because each zone implies a level of trust, the implicit rule for passing traffic between two different zones is deny, and the traffic within a zone is permitted. To allow traffic between two different zones, you must create a security rule that allows traffic to flow between them.

While setting up the basic framework for securing the enterprise perimeter, it's good idea to start with a simple security policy that allows traffic between the different zones without being too restrictive. As illustrated in the following section, our objective is to minimize the likelihood of breaking applications that users on the network need access to, while providing visibility into the applications and the potential threats for your network.

When defining policies make sure that you do not create a policy that denies all traffic from *any* source zone to *any* destination zone as this will break intra-zone traffic that is implicitly allowed. By default, intra-zone traffic is permitted because the source and destination zones are the same and therefore share the same level of trust.

## Define Basic Security Rules

**Step 1** Permit Internet access for all users on the enterprise network.

Zone: Trust to Untrust

By default, the firewall includes a security rule named *rule1* that allows all traffic from Trust zone to Untrust zone. You can either delete the rule or modify the rule to reflect your zone-naming convention.

To safely enable applications that are required for day-to-day business operations we will create a simple rule that allows access to the Internet. To provide basic threat protection, we will attach the default security profiles available on the firewall.

1. Select **Policies > Security** and click **Add.**
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, Set the **Destination Zone** to Untrust.

To scan policy rules and visually identify the zones on each rule, create a tag with the same name as the zone. For example, to color code the Trust zone as green, select **Objects > Tags,** click **Add** and **Name** the tag Trust, and select the **Color** green.

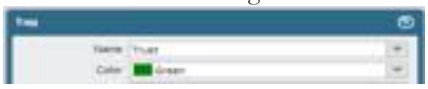

5. In the **Service/ URL Category** tab, select service-http and service-https.
6. In the **Actions** tab, complete these tasks:
   a. Set the **Action Setting** to **Allow**.
   b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection and URL filtering, under **Profile Setting.**
7. Verify that logging is enabled at the end of a session under **Options**. Only traffic that matches a security rule will be logged.

**Step 2** Permit users on the internal network to access the servers in the DMZ.

Zone: Trust to DMZ

If using IP addresses for configuring access to the servers in the DMZ, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT addresses), and the post-NAT zone.

1. Click **Add** in the **Policies > Security** section.
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, set the **Destination Zone** to DMZ.
5. In the **Service/ URL Category** tab, make sure the **Service** is set to **application-default**.
6. In the **Actions** tab, set the **Action Setting** to Allow.
7. Leave all the other options at the default values.

© Palo Alto Networks, Inc.

**Define Basic Security Rules (Continued)**

Step 3 Restrict access from the Internet to the servers on the DMZ to specific server IP addresses only.

For example, you might only allow users to access the webmail servers from outside.

Zone: Untrust to DMZ

To restrict inbound access to the DMZ from the Internet, configure a rule that allows access only to specific servers IP addresses and on the default ports that the applications use.

1. Click **Add** to add a new rule, and give it a descriptive name.
2. In the **Source** tab, set the **Source Zone** to Untrust.
3. In the **Destination** tab, set the **Destination Zone** to DMZ.
4. Set the **Destination Address** to the **Public web server** address object you created earlier. The public web server address object references the public IP address—208.80.56.11/24—of the web server that is accessible on the DMZ.
5. Select the webmail application in the **Application** tab.

   The **Service** is set to **application-default** by default.

6. Set the **Action Setting** to **Allow**.

Step 4 Allow access from the DMZ to your internal network (Trust zone). To minimize risk, you will allow traffic only between specific servers and destination addresses. For example, if you have an application server on the DMZ that needs to communicate with a specific database server in your Trust zone, create a rule to allow traffic between a specific source to a specific destination.

Zone: DMZ to Trust

1. Click **Add** to add a new rule, and give it a descriptive name.
2. Set the **Source Zone** to DMZ.
3. Set the **Destination Zone** to Trust.
4. Create a an address object that specifies the server(s) on your Trust zone that can be accessed from the DMZ.

5. In the **Destination** tab on the Security Policy rule, set the **Destination Address** to the Address object you created above.
6. In the **Actions** tab, complete these tasks:
   a. Set the **Action Setting** to **Allow**.
   b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection, under **Profile Setting.**
   c. In the Other Settings section, select the option to **Disable Server Response Inspection**. This setting disables the antivirus and anti-spyware scanning on the server-side responses, and thus reduces the load on the firewall.

| Define Basic Security Rules (Continued) | | |
|---|---|---|
| Step 5 | Enable the servers on the DMZ to obtain updates and hot fixes from the Internet. Say, for example, you would like to allow the Microsoft Update service.<br><br>Zone: DMZ to Untrust | 1. Add a new rule and give it a descriptive label.<br>2. Set the **Source Zone** to DMZ.<br>3. Set the **Destination Zone** to Untrust.<br>4. Create an application group to specify the applications that you would like to allow. In this example, we allow Microsoft updates (ms-updates) and dns.<br> <br><br>The **Service** is set to **application-default** by default. This allows the firewall to permit the applications only when they use the standard ports associated with these applications.<br>5. Set the **Action Setting** to **Allow**.<br>6. Attach the default profiles for antivirus, anti-spyware, and vulnerability protection, under **Profiles**.<br> 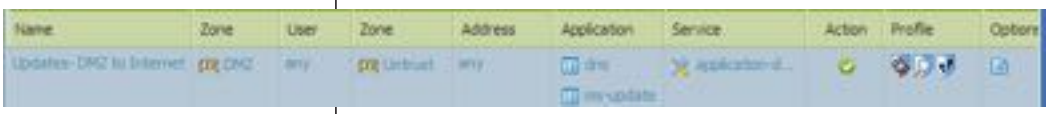 <br> |
| Step 6 | Save your policies to the running configuration on the device. | Click **Commit**. |

## Test Your Security Policies

To verify that you have set up your basic policies effectively, test whether your security policies are being evaluated and determine which security rule applies to a traffic flow.

 © Palo Alto Networks, Inc.

| Verify Policy Match Against a Flow | |
|---|---|
| To verify the policy rule that matches a flow, use the following CLI command:<br>**test security-policy-match source** *<IP_address>* **destination** *<IP_address>* **destination port** *<port_number>* **protocol** *<protocol_number>*<br>The output displays the best rule that matches the source and destination IP address specified in the CLI command. | For example, to verify the policy rule that will be applied for a server on the DMZ with the IP address 208.90.56.11 when it accesses the Microsoft update server, you will try the following command:<br>`test security-policy-match source 208.80.56.11 destination 176.9.45.70 destination-port 80 protocol 6`<br><br>`"Updates-DMZ to Internet" {`<br>`from dmz;`<br>`source any;`<br>`source-region any;`<br>`to untrust;`<br>`destination any;`<br>`destination-region any;`<br>`user any;`<br>`category any;`<br>`application/service[ dns/tcp/any/53 dns/udp/any/53 dns/udp/any/5353 ms-update/tcp/any/80 ms-update/tcp/any/443];`<br>`action allow;`<br>`terminal yes;` |

## Monitor the Traffic on Your Network

Now that you have a basic security policy in place, you can review the statistics and data in the Application Command Center (ACC), traffic logs, and the threat logs to observe trends on your network, to identify where you need to create more granular policies.

Unlike traditional firewalls that use port or protocol to identify applications, the Palo Alto Networks firewalls use the application signature (the App-ID technology) to monitor applications. The application signature is based on unique application properties and related transaction characteristics in combination with the port or protocol. Therefore, even when the traffic uses the right port/protocol, the firewall can deny access to content because the application signature is not a match. This feature allows you to safely enable applications by allowing parts of the application while blocking or controlling functions within the same application. For example, if you allow the application *web-browsing* a user will be able to access content on the Internet. Then, if a user goes to Facebook and then goes on to play Scrabble on Facebook, the firewall will identify the application shifts and recognize Facebook as an *application* and Scrabble as a *Facebook-app*. Therefore, if you create a specific rule that blocks Facebook applications, the user will be denied access to Scrabble while still being able to access Facebook.

To monitor traffic on your network:

- Use the Application Command Center—In the ACC, review the most used applications and the high-risk applications on your network. The ACC graphically summarizes the log information to highlight the applications traversing the network, who is using them (with User-ID enabled), and the potential security

impact of the content to help you identify what is happening on the network in real time. You can then use this information to create appropriate security policies that block unwanted applications, while allowing and enabling applications in a secure manner.

- Determine what updates/modifications are required for your network security rules and implement the changes. For example:
  - Evaluate whether to allow content based on schedule, users, or groups
  - Allow or control certain applications or functions within an application
  - Decrypt and inspect content
  - Allow but scan for threats and exploits

    For information on refining your security policies and for attaching custom security profiles, see Enable Basic Threat Prevention Features.
- View the Log Files—Specifically, view the traffic and threat logs (**Monitor > Logs**).

Traffic logs are dependent on how your security policies are defined and setup to log traffic. The ACC tab, however, records applications and statistics regardless of policy configuration; it shows all traffic that is allowed on your network, therefore it includes the inter zone traffic that is allowed by policy and the same zone traffic that is allowed implicitly.

- Interpret the URL Filtering Logs—Review the URL filtering logs to scan through alerts, denied categories/URL. URL logs are generated when a traffic matches a security rule that has a URL filtering profile attached with an action of alert, continue, override or block.

 © Palo Alto Networks, Inc.

# Enable Basic Threat Prevention Features

The Palo Alto Networks next-generation firewall has unique threat prevention capabilities that allow it to protect your network from attack despite evasive, tunneled, or circumvention techniques. The threat prevention features on the firewall include the WildFire service, the Security Profiles that support Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking and Data Filtering capabilities and the Denial of Service (DoS) and Zone protection functionality.

Before you can apply threat prevention features, you must first configure zones—to identify one or more source or destination interfaces—and security policies. To configure interfaces, zones, and the policies that are needed to apply threat prevention features, see Set Up Interfaces and Zones and Set Up Basic Security Policies.

To begin protecting your network from threats start here:

- Enable WildFire
- Scan Traffic for Threats
- Control Access to Web Content

## Enable WildFire

The WildFire service is included as part of the base product. The WildFire service enables the firewall to forward attachments to a sandbox environment where applications are run to detect any malicious activity. As new malware is detected by the WildFire system, malware signatures are automatically generated and are made available within 24-48 hours in the antivirus daily downloads. Your threat prevention subscription entitles you for antivirus signature updates that include signatures discovered by WildFire.

Consider purchasing the WildFire subscription service for these additional benefits:

- Sub-hourly (as often as every 15 minutes) WildFire signature updates
- Advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet)
- Ability to upload files using the WildFire API
- Ability to forward files to a private WF-500 WildFire appliance

  While the ability to configure a file blocking profile to forward Portable Executable (PE) files to the WildFire cloud for analysis is free, in order to forward files to a private WildFire appliance, a WildFire subscription is required.

| **Enable WildFire** | | |
|---|---|---|
| Step 1 | Confirm that your device is registered and that you have a valid support account as well as any subscriptions you require. | 1. Go to the Palo Alto Networks Support Site, log in, and select **My Devices**.<br>2. Verify that the firewall is listed. If it is not listed, see Register With Palo Alto Networks.<br>3. (Optional) Activate Licenses. |
| Step 2 | Set the WildFire forwarding options.<br><br>If you do not have a WildFire subscription you can only forward executables. | 1. Select **Device > Setup > WildFire**.<br>2. Click the edit icon in the General Settings section.<br>3. (Optional) Specify the **WildFire Server** to which to forward files. By default, the firewall will forward files to the public WildFire cloud hosted in the United States. To forward files to a different WildFire cloud, enter a new value as follows:<br>• To forward to a private WildFire cloud, enter the IP address or FQDN of your WF-500 WildFire appliance.<br>• To forward files to the public WildFire cloud running in Japan, enter `wildfire.paloaltonetworks.jp`.<br>4. (Optional) If you want to change the maximum file size that the firewall can forward for a specific type of file, modify the value in the corresponding field.<br>5. Click **OK** to save your changes. |

 © Palo Alto Networks, Inc.

| Enable WildFire (Continued) | | |
|---|---|---|
| Step 3 | Set up a file blocking profile to forward files to WildFire. | 1. Select **Objects > Security Profiles > File Blocking** and click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the profile.<br>3. Click **Add** to create a forwarding rule and enter a name.<br>4. In the **Action** column, select **forward**.<br>5. Leave the other fields set to **any** to forward any supported file type from any application.<br>6. Click **OK** to save the profile. |
| Step 4 | Attach the file blocking profile to the security policies that allow access to the Internet. | 1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Create Security Rules.<br>2. Click the **Actions** tab within the security policy.<br>3. In the Profile Settings section, click the drop-down and select the file blocking profile you created for WildFire forwarding. (If you don't see a drop-down for selecting a profile, select **Profiles** from the **Profile Type** drop-down. |
| Step 5 | Save the configuration. | Click **Commit**. |
| Step 6 | Verify that the firewall is forwarding files to WildFire. | 1. Select **Monitor > Logs > Data Filtering**.<br>2. Check the **Action** column for the following actions:<br>• **Forward**— Indicates that the file was successfully forwarded by the file blocking profile attached to the security policy.<br>• **Wildfire-upload-success**—Indicates that the file was sent to WildFire. This means the file is not signed by a trusted file signer and it has not been previously analyzed by WildFire.<br>• **Wildfire-upload-skip**—Indicates that the file was identified as eligible to be sent to WildFire by a file blocking profile/security policy, but did not need to be analyzed by WildFire because it has already been analyzed previously. In this case, the action will display as **forward** in the Data Filtering log because it was a valid forward action, but it was not sent to WildFire and analyzed because the file has already been sent to the WildFire cloud from another session, possibly from another firewall.<br>3. View the WildFire logs by selecting **Monitor > Logs > WildFire Submissions**. If new WildFire logs appear, the firewall is successfully forwarding files to WildFire and WildFire is returning file analysis reports. |

# Scan Traffic for Threats

Security profiles provide threat protection in security policies. For example, you can apply an antivirus profile to a security policy and all traffic that matches the security policy will be scanned for viruses.

The following sections provide steps for setting up a basic threat prevention configuration:

- Set Up Antivirus, Anti-Spyware, and Vulnerability Protection
- Set Up File Blocking

## Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

Every Palo Alto Networks next-generation firewall comes with predefined Antivirus, Anti-Spyware, and Vulnerability Protection profiles that you can attach to security policies. There is one predefined Antivirus profile, **default**, which uses the default action for each protocol (block HTTP, FTP, and SMB traffic and alert on SMTP, IMAP, and POP3 traffic). There are two predefined Anti-spyware and Zone Protection profiles:

- **default**—Applies the default action to all client and server critical, high, and medium severity spyware/vulnerability protection events. It does not detect low and informational events.
- **strict**—Applies the block response to all client and server critical, high and medium severity spyware/vulnerability protection events and uses the default action for low and informational events.

To ensure that the traffic entering your network is free from threats, attach the predefined profiles to your basic web access policies. As you monitor the traffic on your network and expand your policy rulebase, you can then design more granular profiles to address your specific security needs.

| **Set up Antivirus/Anti-Spyware/Vulnerability Protection** | | |
|---|---|---|
| Step 1 | Verify that you have a Threat Prevention license. | • The Threat Prevention license bundles the Antivirus, Anti-Spyware, and the Vulnerability Protection features in one license.<br>• Select **Device > Licenses** to verify that the **Threat Prevention** license is installed and valid (check the expiration date). |
| Step 2 | Download the latest antivirus threat signatures. | 1. Select **Device > Dynamic Updates** and click **Check Now** at the bottom of the page to retrieve the latest signatures.<br>2. In the **Actions** column, click **Download** to install the latest Antivirus, and Applications and Threats signatures. |

 © Palo Alto Networks, Inc.

**Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued)**

Step 3 Schedule signature updates.

Perform a **download-and-install** on a daily basis for antivirus updates and weekly for applications and threats updates.

1. From **Device > Dynamic Updates**, click the text to the right of **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.
2. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select **Download Only**, you would need to manually go in and click the **Install** link in the **Action** column to install the signature. When you click **OK**, the update is scheduled. No commit is required.
3. (Optional) You can also enter the number of hours in the **Threshold** field to indicate the minimum age of a signature before a download will occur. For example, if you entered **10**, the signature must be at least 10 hours old before it will be downloaded, regardless of the schedule.
4. In an HA configuration, you can also click the **Sync To Peer** option to synchronize the content update with the HA peer after download/install. This will not push the schedule settings to the peer device, you need to configure the schedule on each device.

Recommendations for HA Configurations:

- **Active/Passive HA**—If the MGT port is used for antivirus signature downloads, you should configure a schedule on both devices and both devices will download/install independently. If you are using a data port for downloads, the passive device will not perform downloads while it is in the passive state. In this case you would set a schedule on both devices and then select the **Sync To Peer** option. This will ensure that whichever device is active, the updates will occur and will then push to the passive device.
- **Active/Active HA**—If the MGT port is used for antivirus signature downloads on both devices, then schedule the download/install on both devices, but do not select the **Sync To Peer** option. If you are using a data port, schedule the signature downloads on both devices and select **Sync To Peer**. This will ensure that if one device in the active/active configuration goes into the active-secondary state, the active device will download/install the signature and will then push it to the active-secondary device.

| Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued) | |
|---|---|
| Step 4 Attach the security profiles to a security policy.<br><br>Attach a clone of a predefined security profile to your basic security policies. That way, if you want to customize the profile you can do so without deleting the read-only predefined **strict** or **default** profile and attaching a customized profile. | 1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab.<br>2. In **Profile Settings**, click the drop-down next to each security profile you would like to enable. In this example we choose default for **Antivirus, Vulnerability Protection,** and **Anti-Spyware**.<br>(If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.)<br> 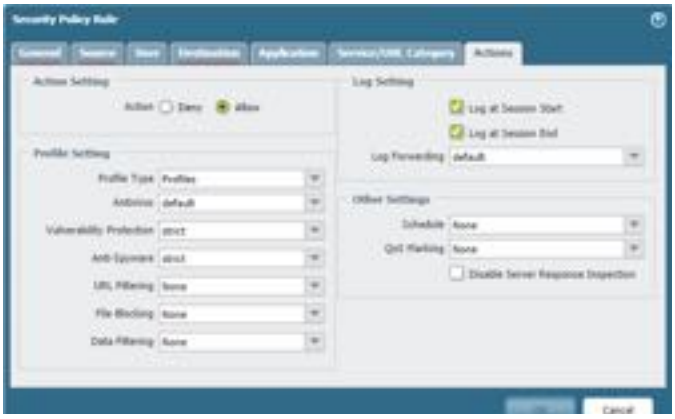 |
| Step 5 Save the configuration. | Click **Commit**. |

## Set Up File Blocking

File blocking profiles allow you to identify specific file types that you want to want to block or monitor. The following workflow shows how to set up a basic file blocking profile that prevents users from downloading executable files from the Internet.

| Configure File Blocking | |
|---|---|
| Step 1 Create the file blocking profile. | 1. Select **Objects > Security Profiles > File Blocking** and click **Add**.<br>2. Enter a **Name** for the file blocking profile, for example *Block_EXE*. Optionally enter a **Description**, such as *Block users from downloading exe files from websites.* |

 © Palo Alto Networks, Inc.

| Configure File Blocking (Continued) | | |
|---|---|---|
| Step 2 | Configure the file blocking options. | 1. Click **Add** to define the profile settings.<br>2. Enter a **Name**, such as *BlockEXE*.<br>3. Set the **Applications** to which to apply file blocking, or leave it set to any.<br>4. Set **File Types** to block. For example, to block download of executables, you would select **exe**.<br>5. Specify the **Direction** in which to block files **download**, **upload**, or **both**.<br>6. Set the **Action** to one of the following:<br>• **continue**—Users will have to click Continue in order to proceed with the download/upload. You must enable response pages on the associated interfaces if you plan to use this option.<br>• **block**—Files matching the selected criteria will be blocked from download/upload.<br>• **alert**—Files matching the selected criteria will be allowed, but will generate a log entry in the data filtering log.<br> <br>7. Click **OK** to save the profile. |
| Step 3 | Attach the file blocking profile to the security policies that allow access to content. | 1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Create Security Rules.<br>2. Click the **Actions** tab within the security policy.<br>3. In the Profile Settings section, click the drop-down and select the file blocking profile you created. (If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.) |
| Step 4 | Enable Response Pages in the management profile for each interface on which you are attaching file blocking profile with a **continue** action. | 1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.<br>2. Select **Response Pages**, as well as any other management services required on the interface.<br>3. Click **OK** to save the interface management profile.<br>4. Select **Network > Interfaces** and select the interface to which to attach the profile.<br>5. On the **Advanced > Other Info** tab, select the interface management profile you just created.<br>6. Click **OK** to save the interface settings. |

**Configure File Blocking (Continued)**

Step 5 To test the file blocking configuration, access a client PC in the trust zone of the firewall and attempt to download an.exe file from a website in the untrust zone. A response page should display. Click **Continue** to download the file. You can also set other actions, such as alert only, forward (which will forward to WildFire), or block, which will not provide a continue page to the user. The following shows the default response page for File Blocking:

**Example: Default File Blocking Response Page**

 © Palo Alto Networks, Inc.

## Control Access to Web Content

URL Filtering provides visibility and control over web traffic on your network. With URL filtering enabled, the firewall can categorize web traffic into one or more (from approximately 60) categories. You can then create policies that specify whether to allow, block, or log (alert) traffic based on the category to which it belongs. The following workflow shows how to enable PAN-DB for URL filtering, create security profiles, and attach them to security policies to enforce a basic URL filtering policy.

| **Configure URL Filtering** | |
|---|---|
| Step 1 Confirm license information for URL Filtering. | 1. Obtain and install a URL Filtering license. See Activate Licenses for details.<br>2. Select **Device > Licenses** and verify that the URL Filtering license is valid.<br> |
| Step 2 Download the seed database and activate the license. | 1. To download the seed database, click **Download** next to **Download Status** in the PAN-DB URL Filtering section of the Licenses page.<br>2. Choose a region (North America, Europe, APAC, Japan) and then click **OK** to start the download.<br>3. After the download completes, click **Activate**.<br> |
| Step 3 Create a URL filtering profile.<br>Best Practice for New Profiles:<br>Because the default URL filtering profile blocks risky and threat-prone content, clone this profile when creating a new profile in order to preserve the default settings. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select the default profile and then click **Clone**. The new profile will be named *default-1*.<br>3. Select the new profile and rename it. |

| Configure URL Filtering (Continued) | |
|---|---|
| Step 4 Define how to control access to web content.<br><br>If you are not sure what traffic you want to control, consider setting the categories (except for those blocked by default) to alert. You can then use the visibility tools on the firewall, such as the ACC and App Scope, to determine which web categories to restrict to specific groups or to block entirely. You can then go back and modify the profile to block and allow categories as desired.<br><br>You can also define specific sites to always allow or always block regardless of category and enable the safe search option to filter search results when defining the URL Filtering profile. | 1. For each category that you want visibility into or control over, select a value from the **Action** column as follows:<br>• If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **Allow**.<br>• For visibility into traffic to sites in a category, select **Alert**<br>• To prevent access to traffic that matches the associated policy, select **Block** (this also generates a log entry).<br> 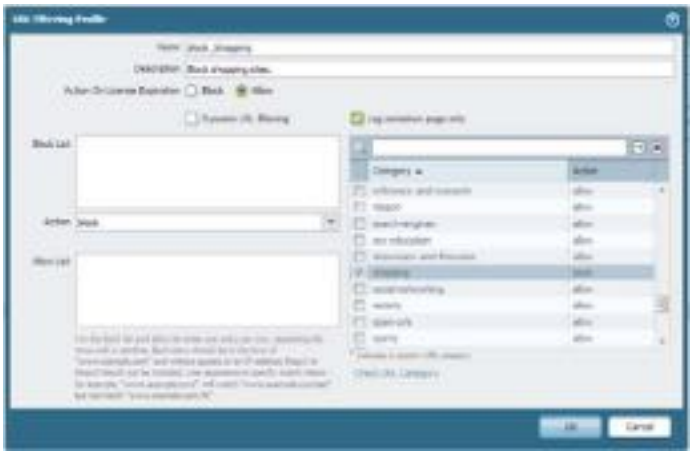 <br>2. Click **OK** to save the URL filtering profile. |
| Step 5 Attach the URL filtering profile to a security policy. | 1. Select **Policies > Security**.<br>2. Select the desired policy to modify it and then click the **Actions** tab.<br>3. If this is the first time you are defining a security profile, select **Profiles** from the **Profile Type** drop-down.<br>4. In the **Profile Settings** list, select the profile you just created from the **URL Filtering** drop-down. (If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.)<br>5. Click **OK** to save the profile.<br>6. **Commit** the configuration. |
| Step 6 Enable Response Pages in the management profile for each interface on which you are filtering web traffic. | 1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.<br>2. Select **Response Pages**, as well as any other management services required on the interface.<br>3. Click **OK** to save the interface management profile.<br>4. Select **Network > Interfaces** and select the interface to which to attach the profile.<br>5. On the **Advanced > Other Info** tab, select the interface management profile you just created.<br>6. Click **OK** to save the interface settings. |

 © Palo Alto Networks, Inc.

| Configure URL Filtering (Continued) | | |
|---|---|---|
| Step 7 | Save the configuration. | Click **Commit**. |
| Step 8 | To test URL filtering, access a client PC from the zone where the security policy is applied and attempt to access a site in a blocked category. You should see a URL Filtering response page that indicates that the page has been blocked: | |

## For More Information

For more detailed information on how to protect your enterprise from threats, see Threat Prevention. For details on how to scan encrypted (SSH or SSL) traffic for threats, see Decryption.

For information about the threats and applications that Palo Alto Networks products can identify, visit the following links:

- Applipedia—Provides details on the applications that Palo Alto Networks can identify.
- Threat Vault—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

# Best Practices for Completing the Firewall Deployment

Now that you have integrated the firewall into your network and enabled the basic security features, you can begin configuring more advanced features. Here are some things to consider next:

- Learn about the different Management Interfaces that are available to you and how to access and use them.
- Set up High Availability—High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in a two-device cluster provides redundancy and allows you to ensure business continuity.
- Configure the Master Key—Every Palo Alto Networks firewall has a default master key that encrypts private keys that are used to authenticate administrators when they access management interfaces on the firewall. As a best practice to safeguard the keys, configure the master key on each firewall to be unique.
- Manage Firewall Administrators—Every Palo Alto Networks firewall and appliance is preconfigured with a default administrative account (admin) that provides full read-write access (also known as superuser access) to the device. As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This allows you to better protect the device from unauthorized configuration (or modification) and to enable logging of the actions of each individual administrator.
- Enable User Identification (User-ID)—User-ID is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses.
- Enable Decryption—Palo Alto Networks firewalls provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted or tunneled traffic.
- Passive DNS Collection—Enable this opt-in feature to enable the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis in order to improve threat intelligence and threat prevention capabilities.
- Follow the Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions.

 © Palo Alto Networks, Inc.

# Device Management

Administrators can configure, manage, and monitor the Palo Alto Networks' firewalls using the web interface, the CLI, and the API management interfaces. Role-based administrative access to the management interfaces can be customized in order to delegate specific tasks or permissions to certain administrators. See the following topics for information on device management options, including how to begin using the management interfaces and how to customize administrator roles:

- Management Interfaces
- Manage Firewall Administrators
- Reference: Web Interface Administrator Access

# Management Interfaces

PAN-OS firewalls and Panorama provide three user interfaces: a web interface, a command line interface (CLI), and a REST management API. See the following topics for how to access and begin using each of the device management interfaces:

- Use the Web Interface to complete administrative tasks and generate reports from the web interface with relative ease. This graphical interface allows you to access the firewall using HTTPS and it is the best way to perform administrative tasks.
- Use the Command Line Interface (CLI) to type through the commands in rapid succession to complete a series of tasks. The CLI is a no-frills interface that supports two command modes and each mode has its own hierarchy of commands and statements. When you get familiar with the nesting structure and the syntax of the commands, the CLI allows quick response times and offers administrative efficiency.
- Use the XML API to streamline your operations and integrate with existing, internally developed applications and repositories. The XML API is provided as a web service that is implemented using HTTP/HTTPS requests and responses.

 © Palo Alto Networks, Inc.

# Use the Web Interface

The following topics describes how to begin using the firewall web interface:

- Launch the Web Interface
- Navigate the Web Interface
- Commit Changes
- Use Configuration Pages
- Required Fields
- Lock Transactions

## Launch the Web Interface

The following web browsers are supported for access to the web interface for PAN-OS firewalls and Panorama:

- Internet Explorer 7+
- Firefox 3.6+
- Safari 5+
- Chrome 11+

Launch an Internet browser and enter the firewall's IP address. Enter your user credentials. If logging in to the firewall for the first time, type the default **admin** into both the **Name** and **Password** fields.

To view information on how to use a specific page and an explanation of the fields and options on the page, click the **Help** icon

in the upper right area of the page to open the online help system. In addition to displaying context-sensitive help for a page, clicking the **Help** icon displays a help navigation pane with options to browse and search all help content.

## Navigate the Web Interface

The following conventions apply when using the web interface.

- To display the menu items for a general functional category, click the tab, such as **Objects** or **Device**, near the top of the browser window.

- Click an item on the side menu to display a panel.

- To display submenu items, click the ⊞ icon to the left of an item. To hide submenu items, click the ⊟ icon to the left of the item.

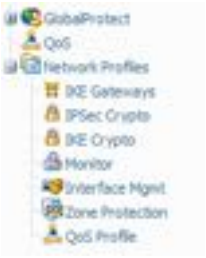


- On most configuration pages, you can click **Add** to create a new item.

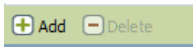


- To delete one or more items, select their check boxes and click **Delete**. In most cases, the system prompts you to confirm by clicking **OK** or to cancel the deletion by clicking **Cancel**.

- On some configuration pages, you can select the check box for an item and click **Clone** to create a new item with the same information as the selected item.

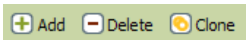


- To modify an item, click its underlined link.

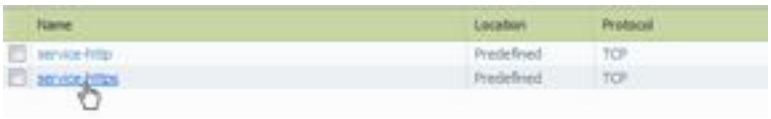


- To view the current list of tasks, click the **Tasks** icon in the lower right corner of the page. The Task Manager window opens to show the list of tasks, along with status, start times, associated messages, and actions. Use the **Show** drop-down list to filter the list of tasks.

 © Palo Alto Networks, Inc.

- The web interface language is controlled by the current language of the computer that is managing the device if a specific language preference has not been defined. For example, if the computer you use to manage the firewall has a locale of Spanish, when you log in to the firewall, the web interface will be in Spanish.
- To specify a language that will always be used for a given account regardless of the locale of the computer, click the **Language** icon in the lower right corner of the page and the Language Preference window opens. Click the drop-down list to select the desired language and then click **OK** to save your change.

- On pages that list information you can modify (for example, the **Setup** page on the **Devices** tab), click the icon in the upper right corner of a section to edit the settings.

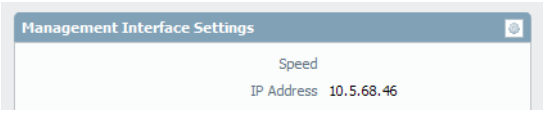


- After you configure settings, you must click **OK** or **Save** to store the changes. When you click **OK**, the current "candidate" configuration is updated.

## Commit Changes

Click **Commit** at the top of the web interface to open the Commit dialog box, which displays the following options. Certain options only display if you click **Advanced**.

- **Include Device and Network configuration**—Include the device and network configuration changes in the commit operation.
- **Include Shared Object configuration**—(Multi-virtual system firewalls only) Include the shared object configuration changes in the commit operation.

- **Include Policy and Objects**—(Non-multi-virtual system firewalls only) Include the policy and object configuration changes in the commit operation.
- **Include virtual system configuration**—Include all virtual systems or choose **Select one or more virtual systems**.
- **Preview Changes**—Click this button to bring up a two-pane window that shows proposed changes in the candidate configuration compared to the current running configuration. You can choose the number of lines of context to display, or show all lines. Changes are color coded based on items that you and other administrators added (green), modified (yellow), or deleted (red) since the last commit.

Because the preview results display in a new window, your browser must allow pop-ups. If the preview window does not open, refer to your browser documentation for the steps to unblock pop-ups.

## Use Configuration Pages

The tables on configuration pages include sorting and column chooser options. Click a column header to sort on that column, and click again to change the sort order. Click the arrow to the right of any column and select check boxes to choose the columns to display.

## Required Fields

Required fields are shown with a light yellow background. A message indicating that the field is required appears when you hover over or click in the field entry area.

## Lock Transactions

The web interface provides support for multiple administrators by allowing an administrator to lock a current set of transactions, thereby preventing configuration changes or commit operations by another administrator until the lock is removed. The following types of locks are supported:

- **Config lock**—Blocks other administrators from making changes to the configuration. This type of lock can be set globally or for a virtual system. It can be removed only by the administrator who set it or by a superuser on the system.

 © Palo Alto Networks, Inc.

- **Commit Lock**—Blocks other administrators from committing changes until all of the locks have been released. This type of lock prevents collisions that can occur when two administrators are making changes at the same time and the first administrator finishes and commits changes before the second administrator has finished. The lock is released when the current changes are committed by the administrator who applied the lock, or it can be released manually.

Any administrator can open the lock window to view the current transactions that are locked, along with a time stamp for each.

To lock a transaction, click the unlocked icon on the top bar to open the Locks dialog box. Click **Take a Lock**, select the scope of the lock from the drop-down list, and click **OK**. Add additional locks as needed, and then click **Close** to close the Lock dialog box.

The transaction is locked, and the icon on the top bar changes to a locked icon that shows the number of locked items in parentheses.

To unlock a transaction, click the locked icon on the top bar to open the Locks window. Click the icon for the lock that you want to remove, and click **Yes** to confirm. Click **Close** to close the Lock dialog box.

You can arrange to automatically acquire a commit lock by selecting the **Automatically acquire commit lock** check box in the Management area of the **Device Setup** page.

# Use the Command Line Interface (CLI)

The PAN-OS CLI allows you to access Firewall and Panorama devices, view status and configuration information, and modify configurations. Access to the PAN-OS CLI is provided through SSH, Telnet, or direct console access.

The following topics describe how to access and begin using the PAN-OS CLI:

- Access the PAN-OS CLI
- Operational and Configuration Modes

For more information on the CLI, refer to the *PAN-OS Command Line Interface Reference Guide*.

## Access the PAN-OS CLI

Before you begin, verify that the firewall is installed and that a SSH, Telnet, or direct console connection is established.

Use the following settings for direct console connection:

- Data rate: 9600
- Data bits: 8
- Parity: none
- Stop bits: 1
- Flow control: None

| **Access the PAN-OS CLI** | |
|---|---|
| Step 1 | Open the console connection |
| Step 2 | Enter the administrative username. The default is admin. |
| Step 3 | Enter the administrative password. The default is admin. |
| Step 4 | The PAN-OS CLI opens in Operational mode, and the CLI prompt is displayed:<br>username@hostname> |

## Operational and Configuration Modes

When you log in, the PAN-OS CLI opens in Operational mode. You can move between Operational and Configuration modes at any time. Use Operational mode to view the state of the system, navigate the PAN-OS CLI, and enter configuration mode. Use Configuration mode to view and modify the configuration hierarchy.

- To enter Configuration mode from Operational mode, use the `configure` command:

```
username@hostname> configure
Entering configuration mode
[edit]
```

 © Palo Alto Networks, Inc.

```
username@hostname#
```

- To leave Configuration mode and return to Operational mode, use the `quit` or `exit` command:

```
username@hostname# quit
Exiting configuration mode
username@hostname>
```

- To enter an Operational mode command while in Configuration mode, use the **run** command, for example:

```
username@hostname# run ping host 1.1.1.2
PING 1.1.1.1 (1.1.1.1) 56(84) bytes of data
...
username@hostname#
```

- To direct an Operational mode command to a particular VSYS, specify the target VSYS with the following command:

```
username@hostname# set system setting target-vsys <vsys_name>
```

# Use the XML API

Palo Alto Networks XML API uses standard HTTP requests to send and receive data, allowing access to several types of data on the device so the data can be easily integrated with and used in other systems. Use the REST Management API to view a firewall or Panorama's configuration, extract report data in XML format, and execute operational commands. API calls can be made directly from command line utilities such as cURL or wget, or using any scripting or application framework that supports RESTful services. When using the API with command lines tools, both HTTP GET and POST methods are supported.

You must generate an API key in order to being using the XML API. The API key authenticates the user to the firewall, application, or Panorama. After you have generated an API key, you can use the key to perform device configuration and operational tasks, retrieve reports and logs, and import and export files. See Generate an API Key for steps to generate an API key.

The following table shows the URL structure for API requests:

| XML API URL Structure | |
|---|---|
| Prior to PAN-OS 4.1.0 | http(s)://hostname/esp/restapi.esp?request-parameters-values |
| PAN-OS 4.1.0 and later | http(s)://hostname/api/?request-parameters-values |

URL structure item definitions:

- **hostname**—Device's IP address or Domain name.
- **request-parameters-values**—A series of multiple 'parameter=value' pairs separated by the ampersand character (&). These values can either be keywords or data-values in standard or XML format (response data is always in XML format).

There are APIs for PAN-OS, User-ID, and WildFire products. For more information on how to use the API interface, refer to the PAN-OS XML API Usage Guide. To access the online community for developing scripts, visit: https://live.paloaltonetworks.com/community/devcenter.

## Generate an API Key

In order to use the API to manage a firewall or application, an API key is required to authenticate all API calls. Admin account credentials are used to generate API keys.

As a best practice, create a separate admin account for XML-based administration.

| Generate an API key | |
|---|---|
| Step 1 Create an administrator account. | 1. In the web interface, on the **Device > Administrators** tab, click **Add.**<br>2. Enter a login **Name** for the admin.<br>3. Enter and confirm a **Password** for the admin.<br>4. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

| Generate an API key (Continued) | | |
|---|---|---|
| Step 2 | Request an API key. | Replace the hostname, username and password parameters in the following URL with the appropriate values from your administrator account credentials:<br>http(s)://hostname/api/?type=keygen&user=username&password=password<br>The API key is displayed in an XML block. For example:<br><response status="success"><br><result><br><key>0RgWc42Oi0vDx2WRUIUM6A</key><br></result><br></response> |
| Step 3 | (Optional) Revoke or change an API key.<br>For PAN-OS 4.1.0 and later releases, generating an API key using the same administrator account credentials returns unique API keys every time, and all of the keys are valid.<br>You can choose to revoke and then change an API key associated with an administrator account by changing the password associated with the administrator account. Any API keys that were generated using the previous credentials would no longer be valid. | 1. On the **Device > Adminstrators** tab, open the administrator account associated with the API key.<br>2. Enter and confirm a new **Password** for the administrator account.<br>3. Click **OK** and **Commit**.<br>Any API keys associated with the admin account prior to the password change are revoked upon Commit.<br>4. (Optional) Use the updated administrator account credentials to generate a new API key. See Step 2. |

Example work flow using an API key:

Request an API key by entering the URL with the appropriate values in a web browser:

```
https://10.xx.10.50/esp/restapi.esp?type=keygen&user=admin&password=admin
```

Entering the URL displays an XML block that contains the API key:

```
<response status="success">
  <result>
    <key>0RgWc42Oi0vDx2WRUIUM6A</key>
  </result>
</response>
```

Continue to use the API key to create API requests. For example, to generate a report:

```
https://10.xx.10.50/esp/restapi.esp?type=report&reporttype=dynamic&reportname
=top-app-summary&period=last-hour&topn=5&key=0RgWc42Oi0vDx2WRUIUM6A=
```

# Manage Firewall Administrators

Every Palo Alto Networks firewall and appliance is preconfigured with a default administrative account (admin) that provides full read-write access (also known as superuser access) to the device.

As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This allows you to better protect the device from unauthorized configuration (or modification) and to enable logging of the actions of each individual administrator.

The following topics describe the various ways you can set up administrative accounts and provides procedures for setting up basic administrative access:

- Administrative Roles
- Administrative Authentication
- Create an Administrative Account

## Administrative Roles

The way you configure administrator accounts depends on the security requirements within your organization, whether you have existing authentication services you want to integrate with, and how many different administrative roles you require. A *role* defines the type of access the associated administrator has to the system. There are two types of roles you can assign:

- **Dynamic Roles**—These are built-in roles that provide access to the firewall. When new features are added, the firewall automatically updates the definitions of dynamic roles; you never need to manually update them. The following table lists the access privileges associated with dynamic roles.

| Dynamic Role | Privileges |
|---|---|
| Superuser | Full access to the firewall, including defining new administrator accounts and virtual systems. You must have superuser privileges to create an administrative user with superuser privileges. |
| Superuser (read-only) | Read-only access to the firewall. |
| Virtual system administrator | Full access to a selected virtual system (vsys) on the firewall. |
| Virtual system administrator (read-only) | Read-only access to a selected vsys on the firewall. |
| Device administrator | Full access to all firewall settings except for defining new accounts or virtual systems. |
| Device administrator (read-only) | Read-only access to all firewall settings except password profiles (no access) and administrator accounts (only the logged in account is visible). |

- **Admin Role Profiles**—Allow you to create your own role definitions in order to provide more granular access control to the various functional areas of the web interface, CLI and/or XML API. For example, you could create an Admin Role Profile for your operations staff that provides access to the device and network

 © Palo Alto Networks, Inc.

configuration areas of the web interface and a separate profile for your security administrators that provides access to security policy definition, logs, and reports. Keep in mind that with Admin Role Profiles you must update the profiles to explicitly assign privileges for new features/components that are added to the product.

## Administrative Authentication

There are four ways you can authenticate administrative users:

- **Local administrator account with local authentication**—Both the administrator account credentials and the authentication mechanisms are local to the firewall. You can further secure the local administrator account by creating a password profile that defines a validity period for passwords and by setting device-wide password complexity settings.
- **Local administrator account with SSL-based authentication**—With this option, you create the administrator accounts on the firewall, but authentication is based on SSH certificates (for CLI access) or client certificates/common access cards (for the web interface). Refer to the article How to Configure Certificate-based Authentication for the WebUI for details on how to configure this type of administrative access.
- **Local administrator account with external authentication**—The administrator accounts are managed on the local firewall, but the authentication functions are offloaded to an existing LDAP, Kerberos, or RADIUS service. To configure this type of account, you must first create an authentication profile that defines how to access the external authentication service and then create an account for each administrator that references the profile.
- **External administrator account and authentication**—Account administration and authentication are handled by an external RADIUS server. To use this option, you must define Vendor Specific Attributes (VSAs) on your RADIUS server that map to the admin role and, optionally, the virtual system objects you have defined on the Palo Alto Networks device. Refer to the Radius Vendor Specific Attributes (VSA) article for details on how to configure this type of administrative access.

## Create an Administrative Account

Create administrative accounts to define access and administrative privileges for firewall administrators. Because it is common to delegate specific administrative tasks to specific administrators with varying roles, Palo Alto Networks recommends that you create admin role profiles that allow administrators access only to the areas of the management interface that are required to perform their jobs. You can assign the various roles you create to individual administrator accounts and specify access privileges to each management interface: the web interface, the Command Line Interface (CLI), and the REST Management API. By creating admin roles with very granular access privileges, you can ensure that sensitive company data is protected and end user privacy is ensured.

The following procedure describes how to create a local administrator account with local authentication, including how to set administrator access for each management interface.

| Create a Local Administrator | |
|---|---|
| Step 1 Create the Admin Role Profiles that you plan to assign to your administrators (this does not apply if you plan to use Dynamic Roles). The Admin Role Profiles define what type of access to give to the different sections of the web interface, CLI, and XML API for each administrator that you assign a role to.<br><br>You can use this step to set particularly granular access privileges for web interface users. For details on what a specific option enables on the **Web UI** tab, see Web Interface Access Privileges. | Complete the following steps for each role you want to create:<br>1. Select **Device > Admin Roles** and then click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the role.<br>3. On the Web UI, Command Line and/or XML API tabs, specify the access to allow for each management interface:<br>• On the **Web UI** and/or **XML API** tabs, set the access levels for each functional area of the interface by clicking the icon to toggle it to the desired setting: Enable, Read Only, or Disable.<br>• On the **Command Line** tab, specify the type of access to allow to the CLI: **superreader**, **deviceadmin**, or **devicereader** (for Device roles); **vsysadmin** or **vsysreader** (for Virtual System roles); or **None** to disable CLI access entirely.<br>4. Click **OK** to save the profile.<br><br>For example, allow an admin full access to a device using the XML API, with the exception of importing or exporting files:<br> |

 © Palo Alto Networks, Inc.

| Create a Local Administrator (Continued) | | |
|---|---|---|
| Step 2 | (Optional) Set requirements for local user-defined passwords. | • **Create Password Profiles**—Define how often administrators must change their passwords. You can create multiple password profiles and apply them to administrator accounts as needed to enforce the desired security. To create a password profile, select **Device > Password Profiles** and then click the **Add**.<br>• **Configure minimum password complexity settings**—Define rules that govern password complexity, allowing you to force administrators to create passwords that are harder to guess, crack, or compromise. Unlike password profiles, which can be applied to individual accounts, these rules are device wide and apply to all passwords. To configure the settings, select **Device > Setup** and edit the Minimum Password Complexity section. |
| Step 3 | Create an account for each administrator. | 1. Select **Device > Administrators** and then click **Add**.<br>2. Enter a user **Name** and **Password** for the administrator, or create an **Authentication Profile** to use for validating an administrative user's credentials to an external authentication server. See Step 4 for details on setting up an authentication profile.<br>3. Select the **Role** to assign to this administrator. You can either select one of the predefined Dynamic roles or a custom Role Based profile if you created one in Step 1.<br>4. (Optional) Select a **Password Profile**.<br>5. Click **OK** to save the account. |

| Create a Local Administrator (Continued) | | |
|---|---|---|
| Step 4 | (Optional) Set up authentication to an external server—LDAP, RADIUS or Kerberos.<br><br>The server profile specifies how the firewall can connect to the authentication service you plan to use. | 1. Select **Device > Authentication Profile** and then click **Add**.<br>2. Enter a user **Name** to identify the authentication profile.<br>3. Define the conditions for locking out the administrative user.<br>a. Enter the **Lockout Time**. This is the number of minutes that a user is locked out upon reaching the maximum number of failed attempts (0-60 minutes; default 0). 0 means that the lockout is in effect until it is manually unlocked.<br>b. Enter the **Failed Attempts** count. This is the number of failed login attempts that are allowed before the account is locked out (1-10; default 0). By default, the failed attempt count is 0 and the user is not locked out despite repeated failure to authenticate.<br>4. Specify the users and groups that are explicitly allowed to authenticate. By adding an **Allow List** to an authentication profile, you can limit access to specific users in a user group/directory.<br>• Select the **All** check box to allow all users.<br>• Click **Add** and enter the first few characters of a name in the field to list all the users and user groups that start with those characters. Repeat to add as many users/user groups as required.<br>5. In the **Authentication** drop-down, select the type of authentication you plan to use on your network.<br>If you plan to use local database authentication, you must create the local database. Select **Device > Local User Database** and add the users and groups to be authenticated.<br>6. For access to an external authentication server (not local database), select the appropriate server profile in the **Server Profile** drop-down. To create a new server profile, click the link next to New and continue with configuring access to the **LDAP**, **RADIUS** or **Kerberos** server.<br>7. Click **OK**. |
| Step 5 | Commit your changes. | 1. Click **Commit**. |


© Palo Alto Networks, Inc.

# Reference: Web Interface Administrator Access

See the following topics for details on the options to set particularly granular access privileges for PAN-OS and Panorama web interface administrators.

- Web Interface Access Privileges
- Panorama Web Interface Access

## Web Interface Access Privileges

If you want to prevent a role-based administrator from accessing specific tabs on the web interface, you can disable the tab and the administrator will not even see it when logging in using the associated role-based administrative account. For example, you could create an Admin Role Profile for your operations staff that provides access to the **Device** and **Network** tabs only and a separate profile for your security administrators that provides access to the **Object**, **Policy**, and **Monitor** tabs.

The following table describes the tab-level access privileges you can assign to the admin role profile. It also provides cross-references to additional tables that detail granular privileges within a tab. For specific information on how to set the admin role profile to protect end user privacy, see Define User Privacy Settings in the Admin Role Profile.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Dashboard | Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets. | Yes | No | Yes |
| ACC | Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Yes | No | Yes |
| Monitor | Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the admin can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab. | Yes | No | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Policies | Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the admin can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab. | Yes | No | Yes |
| Objects | Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab. | Yes | No | Yes |
| Network | Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the admin can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab. | Yes | No | Yes |
| Device | Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any device-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab.<br> You cannot enable access to the **Admin Roles** or **Administrators** nodes for a role-based administrator even if you enable full access to the **Device** tab. | Yes | No | Yes |


© Palo Alto Networks, Inc.

## Provide Granular Access to the Monitor Tab

In some cases you might want to enable the administrator to view some but not all areas of the Monitor tab. For example, you may want to restrict operations admins to the configuration and system logs only, because they do not contain sensitive user data. Although this section of the admin role definition specifies what areas of the Monitor tab the admin can see, you can also couple privileges in this section with privacy privileges, such as disabling the ability to see user names in logs and reports. One thing to keep in mind, however, is that any system generated reports will still show user names and IP addresses even if you disable that functionality in the role. For this reason, if you do not want the admin to see any of the private user information, you should disable access to the specific reports as detailed in the following table.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Monitor | Enables or disables access to the **Monitor** tab. If disabled, the admin will not see this tab or any of the associated logs or reports. | Yes | No | Yes |
| Logs | Enables or disables access to all log files. You can also leave this privilege enabled and then disable specific logs that you do not want the admin to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the logs, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Yes | No | Yes |
| Traffic | Specifies whether the admin can see the traffic logs. | Yes | No | Yes |
| Threat | Specifies whether the admin can see the threat logs. | Yes | No | Yes |
| URL Filtering | Specifies whether the admin can see the URL filtering logs. | Yes | No | Yes |
| WildFire Submissions | Specifies whether the admin can see the WildFire logs. These logs are only available if you have a WildFire subscription. | Yes | No | Yes |
| Data Filtering | Specifies whether the admin can see the data filtering logs. | Yes | No | Yes |
| HIP Match | Specifies whether the admin can see the HIP Match logs. HIP Match logs are only available if you have a GlobalProtect portal license and gateway subscription. | Yes | No | Yes |
| Configuration | Specifies whether the admin can see the configuration logs. | Yes | No | Yes |
| System | Specifies whether the admin can see the system logs. | Yes | No | Yes |

© Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Alarms | Specifies whether the admin can see system generated alarms. | Yes | No | Yes |
| Packet Capture | Specifies whether the admin can see packet captures (pcaps) from the **Monitor** tab. Keep in mind that packet captures are raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the pcap and you should therefore disable the Packet Capture privilege if you are concerned about user privacy. | Yes | Yes | Yes |
| App Scope | Specifies whether the admin can see the App Scope visibility and analysis tools. Enabling App Scope enables access to all of the **App Scope** charts. | Yes | No | Yes |
| Session Browser | Specifies whether the admin can browse and filter current running sessions on the firewall. Keep in mind that the session browser shows raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the session browser and you should therefore disable the **Session Browser** privilege if you are concerned about user privacy. | Yes | No | Yes |
| Botnet | Specifies whether the admin can generate and view botnet analysis reports or view botnet reports in read-only mode. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in scheduled botnet reports and you should therefore disable the **Botnet** privilege if you are concerned about user privacy. | Yes | Yes | Yes |
| PDF Reports | Enables or disables access to all PDF reports. You can also leave this privilege enabled and then disable specific PDF reports that you do not want the admin to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Yes | No | Yes |
| Manage PDF Summary | Specifies whether the admin can view, add or delete PDF summary report definitions. With read-only access, the admin can see PDF summary report definitions, but not add or delete them. If you disable this option, the admin can neither view the report definitions or add/delete them. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| PDF Summary Reports | Specifies whether the admin can see the generated PDF Summary reports in **Monitor > Reports**. If you disable this option, the **PDF Summary Reports** category will not display in the **Reports** node. | Yes | No | Yes |
| User Activity Report | Specifies whether the admin can view, add or delete User Activity report definitions and download the reports. With read-only access, the admin can see User Activity report definitions, but not add, delete, or download them. If you disable this option, the admin cannot see this category of PDF report. | Yes | Yes | Yes |
| Report Groups | Specifies whether the admin can view, add or delete report group definitions. With read-only access, the admin can see report group definitions, but not add or delete them. If you disable this option, the admin cannot see this category of PDF report. | Yes | Yes | Yes |
| Email Scheduler | Specifies whether the admin can schedule report groups for email. Because the generated reports that get emailed may contain sensitive user data that is not removed by disabling the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** options and because they may also show log data to which the admin does not have access, you should disable the **Email Scheduler** option if you have user privacy requirements. | Yes | Yes | Yes |
| Manage Custom Reports | Enables or disables access to all Custom report functionality. You can also leave this privilege enabled and then disable specific custom report categories that you do not want the admin to be able to access. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option.<br> Reports that are scheduled to run rather than run on demand will show IP address and user information. In this case, be sure to restrict access to the corresponding report areas. In addition, the custom report feature does not restrict the ability to generate reports that contain log data contained in logs that are excluded from the admin role. | Yes | No | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Application Statistics | Specifies whether the admin can create a custom report that includes data from the application statistics database. | Yes | No | Yes |
| Data Filtering Log | Specifies whether the admin can create a custom report that includes data from the data filtering log. | Yes | No | Yes |
| Threat Log | Specifies whether the admin can create a custom report that includes data from the threat log. | Yes | No | Yes |
| Threat Summary | Specifies whether the admin can create a custom report that includes data from the threat summary database. | Yes | No | Yes |
| Traffic Log | Specifies whether the admin can create a custom report that includes data from the traffic log. | Yes | No | Yes |
| Traffic Summary | Specifies whether the admin can create a custom report that includes data from the traffic summary database. | Yes | No | Yes |
| Url Log | Specifies whether the admin can create a custom report that includes data from the URL filtering log. | Yes | No | Yes |
| Hipmatch | Specifies whether the admin can create a custom report that includes data from the HIP match log. | Yes | No | Yes |
| View Scheduled Custom Reports | Specifies whether the admin can view a custom report that has been scheduled to generate. | Yes | No | Yes |
| View Predefined Application Reports | Specifies whether the admin can view Application Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Yes | No | Yes |
| View Predefined Threat Reports | Specifies whether the admin can view Threat Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Yes | No | Yes |
| View Predefined URL Filtering Reports | Specifies whether the admin can view URL Filtering Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| View Predefined Traffic Reports | Specifies whether the admin can view Traffic Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Yes | No | Yes |

## Provide Granular Access to the Policy Tab

If you enable the Policy option in the admin role profile, you can then enable, disable, or provide read-only access to specific nodes within the tab as necessary for the admin role you are defining. By enabling access to a specific policy type, you enable the ability to view, add, or delete policy rules. By enabling read-only access to a specific policy, you enable the admin to view the corresponding policy rule base, but not add or delete rules. Disabling access to a specific type of policy prevents the admin from seeing the policy rule base.

Because policy that is based on specific users (by user name or IP address) must be explicitly defined, privacy settings that disable the ability to see full IP addresses or user names do not apply to the Policy tab. Therefore, you should only allow access to the Policy tab to administrators that are excluded from user privacy restrictions.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Security | Enable this privilege to allow the admin to view, add, and/or delete security policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the security policy rule base, disable this privilege. | Yes | Yes | Yes |
| NAT | Enable this privilege to allow the admin to view, add, and/or delete NAT policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the NAT policy rule base, disable this privilege. | Yes | Yes | Yes |
| QoS | Enable this privilege to allow the admin to view, add, and/or delete QoS policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the QoS policy rule base, disable this privilege. | Yes | Yes | Yes |
| Policy Based Forwarding | Enable this privilege to allow the admin to view, add, and/or delete Policy Based Forwarding (PBF) policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the PBF policy rule base, disable this privilege. | Yes | Yes | Yes |

© Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Decryption | Enable this privilege to allow the admin to view, add, and/or delete decryption policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the decryption policy rule base, disable this privilege. | Yes | Yes | Yes |
| Application Override | Enable this privilege to allow the admin to view, add, and/or delete application override policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the application override policy rule base, disable this privilege. | Yes | Yes | Yes |
| Captive Portal | Enable this privilege to allow the admin to view, add, and/or delete captive portal policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the captive portal policy rule base, disable this privilege. | Yes | Yes | Yes |
| DoS Protection | Enable this privilege to allow the admin to view, add, and/or delete DoS protection policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the DoS protection policy rule base, disable this privilege. | Yes | Yes | Yes |

## Provide Granular Access to the Objects Tab

An *object* is a container that groups specific policy filter values—such as IP addresses, URLs, applications, or services—for simplified rule definition. For example, an address object might contain specific IP address definitions for the web and application servers in your DMZ zone.

When deciding whether to allow access to the objects tab as a whole, determine whether the admin will have policy definition responsibilities. If not, the admin probably does not need access to the tab. If, however, the admin will need to create policy, you can enable access to the tab and then provide granular access privileges at the node level.

By enabling access to a specific node, you give the admin the privilege to view, add, and delete the corresponding object type. Giving read-only access allows the admin to view the already defined objects, but not create or delete any. Disabling a node prevents the admin from seeing the node in the web interface.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Addresses | Specifies whether the admin can view, add, or delete address objects for use in security policy. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Address Groups | Specifies whether the admin can view, add, or delete address group objects for use in security policy. | Yes | Yes | Yes |
| Regions | Specifies whether the admin can view, add, or delete regions objects for use in security, decryption, or DoS policy. | Yes | Yes | Yes |
| Applications | Specifies whether the admin can view, add, or delete application objects for use in policy. | Yes | Yes | Yes |
| Application Groups | Specifies whether the admin can view, add, or delete application group objects for use in policy. | Yes | Yes | Yes |
| Application Filters | Specifies whether the admin can view, add, or delete application filters for simplification of repeated searches. | Yes | Yes | Yes |
| Services | Specifies whether the admin can view, add, or delete service objects for use in creating policies that limit the port numbers an application can use. | Yes | Yes | Yes |
| Service Groups | Specifies whether the admin can view, add, or delete service group objects for use in security policy. | Yes | Yes | Yes |
| Tags (Panorama only) | Specifies whether the admin can view, add, or delete tags that have been defined on the device. | Yes | Yes | Yes |
| GlobalProtect | Specifies whether the admin can view, add, or delete HIP objects and profiles. You can restrict access to both types of objects at the GlobalProtect level, or provide more granular control by enabling the GlobalProtect privilege and restricting HIP Object or HIP Profile access. | Yes | No | Yes |
| HIP Objects | Specifies whether the admin can view, add, or delete HIP objects, which are used to define HIP profiles. HIP Objects also generate HIP Match logs. | Yes | Yes | Yes |
| HIP Profiles | Specifies whether the admin can view, add, or delete HIP Profiles for use in security policy and/or for generating HIP Match logs. | Yes | Yes | Yes |
| Dynamic Block Lists | Specifies whether the admin can view, add, or delete dynamic block lists for use in security policy. | Yes | Yes | Yes |

© Palo Alto Networks, Inc.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Custom Objects | Specifies whether the admin can see the custom spyware and vulnerability signatures. You can restrict access to either enable or disable access to all custom signatures at this level, or provide more granular control by enabling the Custom Objects privilege and then restricting access to each type of signature. | Yes | No | Yes |
| Data Patterns | Specifies whether the admin can view, add, or delete custom data pattern signatures for use in creating custom vulnerability protection profiles. | Yes | Yes | Yes |
| Spyware | Specifies whether the admin can view, add, or delete custom spyware signatures for use in creating custom vulnerability protection profiles. | Yes | Yes | Yes |
| Vulnerability | Specifies whether the admin can view, add, or delete custom vulnerability signatures for use in creating custom vulnerability protection profiles. | Yes | Yes | Yes |
| URL Category | Specifies whether the admin can view, add, or delete custom URL categories for use in policy. | Yes | Yes | Yes |
| Security Profiles | Specifies whether the admin can see security profiles. You can restrict access to either enable or disable access to all security profiles at this level, or provide more granular control by enabling the Security Profiles privilege and then restricting access to each type of profile. | Yes | No | Yes |
| Antivirus | Specifies whether the admin can view, add, or delete antivirus profiles. | Yes | Yes | Yes |
| Anti-Spyware | Specifies whether the admin can view, add, or delete anti-spyware profiles. | Yes | Yes | Yes |
| Vulnerability Protection | Specifies whether the admin can view, add, or delete vulnerability protection profiles. | Yes | Yes | Yes |
| URL Filtering | Specifies whether the admin can view, add, or delete URL filtering profiles. | Yes | Yes | Yes |
| File Blocking | Specifies whether the admin can view, add, or delete file blocking profiles. | Yes | Yes | Yes |
| Data Filtering | Specifies whether the admin can view, add, or delete data filtering profiles. | Yes | Yes | Yes |
| DoS Protection | Specifies whether the admin can view, add, or delete DoS protection profiles. | Yes | Yes | Yes |
| Security Profile Groups | Specifies whether the admin can view, add, or delete security profile groups. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Log Forwarding | Specifies whether the admin can view, add, or delete log forwarding profiles. | Yes | Yes | Yes |
| Decryption Profile | Specifies whether the admin can view, add, or delete decryption profiles. | Yes | Yes | Yes |
| Schedules | Specifies whether the admin can view, add, or delete schedules for limiting a security policy to a specific date and/or time range. | Yes | Yes | Yes |

## Provide Granular Access to the Network Tab

When deciding whether to allow access to the **Network** tab as a whole, determine whether the admin will have network administration responsibilities, including GlobalProtect administration. If not, the admin probably does not need access to the tab.

You can also define access to the **Network** tab at the node level. By enabling access to a specific node, you give the admin the privilege to view, add, and delete the corresponding network configurations. Giving read-only access allows the admin to view the already defined configuration, but not create or delete any. Disabling a node prevents the admin from seeing the node in the web interface.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Interfaces | Specifies whether the admin can view, add, or delete interface configurations. | Yes | Yes | Yes |
| Zones | Specifies whether the admin can view, add, or delete zones. | Yes | Yes | Yes |
| VLANs | Specifies whether the admin can view, add, or delete VLANs. | Yes | Yes | Yes |
| Virtual Wires | Specifies whether the admin can view, add, or delete virtual wires. | Yes | Yes | Yes |
| Virtual Routers | Specifies whether the admin can view, add, modify or delete virtual routers. | Yes | Yes | Yes |
| IPSec Tunnels | Specifies whether the admin can view, add, modify, or delete IPSec Tunnel configurations. | Yes | Yes | Yes |
| DHCP | Specifies whether the admin can view, add, modify, or delete DHCP server and DHCP relay configurations. | Yes | Yes | Yes |
| DNS Proxy | Specifies whether the admin can view, add, modify, or delete DNS proxy configurations. | Yes | Yes | Yes |

© Palo Alto Networks, Inc.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| GlobalProtect | Specifies whether the admin can view, add, modify GlobalProtect portal and gateway configurations. You can disable access to the GlobalProtect functions entirely, or you can enable the GlobalProtect privilege and then restrict the role to either the portal or gateway configuration areas. | Yes | No | Yes |
| Portals | Specifies whether the admin can view, add, modify, or delete GlobalProtect portal configurations. | Yes | Yes | Yes |
| Gateways | Specifies whether the admin can view, add, modify, or delete GlobalProtect gateway configurations. | Yes | Yes | Yes |
| MDM | Specifies whether the admin can view add, modify, or delete GlobalProtect MDM server configurations. | Yes | Yes | Yes |
| QoS | | Yes | Yes | Yes |
| Network Profiles | Sets the default state to enable or disable for all of the Network settings described below. | Yes | No | Yes |
| IKE Gateways | Controls access to the **Network Profiles > IKE Gateways** node. If you disable this privilege, the administrator will not see the **IKE Gateways** node or define gateways that include the configuration information necessary to perform IKE protocol negotiation with peer gateway.<br>If the privilege state is set to read only, you can view the currently configured IKE Gateways but cannot add or edit gateways. | Yes | Yes | Yes |
| IPSec Crypto | Controls access to the **Network Profiles > IPSec Crypto** node. If you disable this privilege, the administrator will not see the **Network Profiles > IPSec Crypto** node or specify protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPSec SA negotiation.<br>If the privilege state is set to read only, you can view the currently configured IPSec Crypto configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| IKE Crypto | Controls how devices exchange information to ensure secure communication. Specify the protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPsec SA negotiation (IKEv1 Phase-1). | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Monitor | Controls access to the **Network Profiles > Monitor** node. If you disable this privilege, the administrator will not see the **Network Profiles > Monitor** node or be able to create or edit a monitor profile that is used to monitor IPSec tunnels and monitor a next-hop device for policy-based forwarding (PBF) rules.<br>If the privilege state is set to read only, you can view the currently configured monitor profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| Interface Mgmt | Controls access to the **Network Profiles > Interface Mgmt** node. If you disable this privilege, the administrator will not see the **Network Profiles > Interface Mgmt** node or be able to specify the protocols that are used to manage the firewall.<br>If the privilege state is set to read only, you can view the currently configured Interface management profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| Zone Protection | Controls access to the **Network Profiles > Zone Protection** node. If you disable this privilege, the administrator will not see the **Network Profiles > Zone Protection** node or be able to configure a profile that determines how the firewall responds to attacks from specified security zones.<br>If the privilege state is set to read only, you can view the currently configured Zone Protection profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| QoS Profile | Controls access to the **Network Profiles > QoS** node. If you disable this privilege, the administrator will not see the **Network Profiles > QoS** node or be able to configure a profile QoS profile that determines how QoS traffic classes are treated.<br>If the privilege state is set to read only, you can view the currently configured QoS profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |

## Provide Granular Access to the Device Tab

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Setup | Controls access to the **Setup** node. If you disable this privilege, the administrator will not see the **Setup** node or have access to device-wide setup configuration information, such as Management, Operations, Service, Content-ID, Wildfire or Session setup information.<br>If the privilege state is set to read only, you can view the current configuration but cannot make any changes. | Yes | Yes | Yes |
| Config Audit | Controls access to the **Config Audit** node. If you disable this privilege, the administrator will not see the **Config Audit** node or have access to any device-wide configuration information. | Yes | No | Yes |
| Admin Roles | Controls access to the **Admin Roles** node. This function can only be allowed for read only access.<br>If you disable this privilege, the administrator will not see the **Admin Roles** node or have access to any device-wide information concerning admin roles configuration.<br>If you set this privilege to read only, you can view the configuration information for all admin roles configured on the device. | No | Yes | Yes |
| Administrators | Controls access to the **Administrators** node. This function can only be allowed for read only access.<br>If you disable this privilege, the administrator will not see the **Administrators** node or have access to information about their own admin account.<br>If you set this privilege to read only, the administrator can view the configuration information for their own admin account. They will not see any information about other admin accounts configured on the device. | No | Yes | Yes |
| Virtual Systems | Controls access to the **Virtual Systems** node. If you disable this privilege, the administrator will not see or be able to configure virtual systems.<br>If the privilege state is set to read only, you can view the currently configured virtual systems but cannot add or edit a configuration. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Shared Gateways | Controls access to the **Shared Gateways** node. Shared gateways allow virtual systems to share a common interface for external communications.<br>If you disable this privilege, the administrator will not see or be able to configure shared gateways.<br>If the privilege state is set to read only, you can view the currently configured shared gateways but cannot add or edit a configuration. | Yes | Yes | Yes |
| User Identification | Controls access to the **User Identification** node. If you disable this privilege, the administrator will not see the **User Identification** node or have access to device-wide User Identification configuration information, such as User Mapping, User-ID Agents, Service, Terminal Services Agents, Group Mappings Settings or Captive Portal Settings.<br>If you set this privilege to read only, the administrator can view configuration information for the device but is not allowed to perform any configuration procedures. | Yes | Yes | Yes |
| VM Information Source | Controls access to the **VM Information Source** node that allows you to configure the firewall/Windows User-ID agent to collect VM inventory automatically. If you disable this privilege, the administrator will not see the **VM Information Source** node.<br>If you set this privilege to read only, the administrator can view the VM information sources configured but cannot add, edit, or delete any sources. | Yes | Yes | Yes |
| High Availability | Controls access to the **High Availability** node. If you disable this privilege, the administrator will not see the **High Availability** node or have access to device-wide high availability configuration information such as General setup information or Link and Path Monitoring.<br>If you set this privilege to read only, the administrator can view High Availability configuration information for the device but is not allowed to perform any configuration procedures. | Yes | Yes | Yes |
| Certificate Management | Sets the default state to enable or disable for all of the Certificate settings described below. | Yes | No | Yes |

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Certificates | Controls access to the **Certificates** node. If you disable this privilege, the administrator will not see the **Certificates** node or be able to configure or access information regarding Device Certificates or Default Trusted Certificate Authorities.<br>If you set this privilege to read only, the administrator can view Certificate configuration information for the device but is not allowed to perform any configuration procedures. | Yes | Yes | Yes |
| Certificate Profile | Controls access to the **Certificate Profile** node. If you disable this privilege, the administrator will not see the **Certificate Profile** node or be able to create certificate profiles.<br>If you set this privilege to read only, the administrator can view Certificate Profiles that are currently configured for the device but is not allowed to create or edit a certificate profile. | Yes | Yes | Yes |
| OCSP Responder | Controls access to the **OCSP Responder** node. If you disable this privilege, the administrator will not see the **OCSP Responder** node or be able to define a server that will be used to verify the revocation status of certificates issues by the PAN-OS device.<br>If you set this privilege to read only, the administrator can view the **OCSP Responder** configuration for the device but is not allowed to create or edit an OCSP responder configuration. | Yes | Yes | Yes |
| Response Pages | Controls access to the **Response Pages** node. If you disable this privilege, the administrator will not see the **Response Page** node or be able to define a custom HTML message that is downloaded and displayed instead of a requested web page or file.<br>If you set this privilege to read only, the administrator can view the **Response Page** configuration for the device but is not allowed to create or edit a response page configuration. | Yes | Yes | Yes |
| Log Settings | Sets the default state to enable or disable for all of the Log settings described below. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| System | Controls access to the **Log Settings > System** node. If you disable this privilege, the administrator will not see the **Log Settings > System** node or be able to specify the severity levels of the system log entries that are logged remotely with Panorama and sent as SNMP traps, syslog messages, and/or email notifications.<br>If you set this privilege to read only, the administrator can view the **Log Settings > System** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |
| Config | Controls access to the **Log Settings > Config** node. If you disable this privilege, the administrator will not see the **Log Settings > Config** node or be able to specify the configuration log entries that are logged remotely with Panorama, and sent as syslog messages and/or email notification.<br>If you set this privilege to read only, the administrator can view the **Log Settings > Config** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |
| HIP Match | Controls access to the **Log Settings > HIP Match** node. If you disable this privilege, the administrator will not see the **Log Settings > HIP Match** node or be able to specify the Host Information Profile (HIP) match log settings that are used to provide information on security policies that apply to GlobalProtect clients<br>If you set this privilege to read only, the administrator can view the **Log Settings > HIP** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |
| Alarms | Controls access to the **Log Settings > Alarms** node. If you disable this privilege, the administrator will not see the **Log Settings > Alarms** node or be able to configure notifications that are generated when a security rule (or group of rules) has been hit repeatedly in a set period of time.<br>If you set this privilege to read only, the administrator can view the **Log Settings > Alarms** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Manage Logs | Controls access to the **Log Settings > Manage Logs** node. If you disable this privilege, the administrator will not see the **Log Settings > Manage Logs** node or be able to clear the indicated logs.<br>If you set this privilege to read only, the administrator can view the **Log Settings > Manage Logs** information but cannot clear any of the logs. | Yes | Yes | Yes |
| Server Profiles | Sets the default state to enable or disable for all of the Server Profiles settings described below. | Yes | No | Yes |
| SNMP Trap | Controls access to the **Server Profiles> SNMP Trap** node. If you disable this privilege, the administrator will not see the **Server Profiles> SNMP Trap** node or be able to specify one or more SNMP trap destinations to be used for system log entries.<br>If you set this privilege to read only, the administrator can view the **Server Profiles> SNMP Trap Logs** information but cannot specify SNMP trap destinations. | Yes | Yes | Yes |
| Syslog | Controls access to the **Server Profiles> Syslog** node. If you disable this privilege, the administrator will not see the **Server Profiles> Syslog** node or be able to specify one or more syslog servers.<br>If you set this privilege to read only, the administrator can view the **Server Profiles> Syslog** information but cannot specify syslog servers. | Yes | Yes | Yes |
| Email | Controls access to the **Server Profiles> Email** node. If you disable this privilege, the administrator will not see the **Server Profiles> Email** node or be able to configure an email profile that can be used to enable email notification for system and configuration log entries<br>If you set this privilege to read only, the administrator can view the **Server Profiles> Email** information but cannot configure and email profile. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Netflow | Controls access to the **Server Profiles> Netflow** node. If you disable this privilege, the administrator will not see the **Server Profiles> Netflow** node or be able to define a NetFlow server profile, which specifies the frequency of the export along with the NetFlow servers that will receive the exported data.<br>If you set this privilege to read only, the administrator can view the **Server Profiles> Netflow** information but cannot define a Netflow profile. | Yes | Yes | Yes |
| RADIUS | Controls access to the **Server Profiles> RADIUS** node. If you disable this privilege, the administrator will not see the **Server Profiles> RADIUS** node or be able to configure settings for the RADIUS servers that are identified in authentication profiles.<br>If you set this privilege to read only, the administrator can view the **Server Profiles> RADIUS** information but cannot configure settings for the RADIUS servers. | Yes | Yes | Yes |
| LDAP | Controls access to the **Server Profiles> LDAP** node. If you disable this privilege, the administrator will not see the **Server Profiles> LDAP** node or be able to configure settings for the LDAP servers to use for authentication by way of authentication profiles.<br>If you set this privilege to read only, the administrator can view the **Server Profiles> LDAP** information but cannot configure settings for the LDAP servers. | Yes | Yes | Yes |
| Kerberos | Controls access to the **Server Profiles> Kerberos** node. If you disable this privilege, the administrator will not see the **Server Profiles> Kerberos** node or configure a Kerberos server that allows users to authenticate natively to a domain controller.<br>If you set this privilege to read only, the administrator can view the **Server Profiles> Kerberos** information but cannot configure settings for Kerberos servers. | Yes | Yes | Yes |
| Local User Database | Sets the default state to enable or disable for all of the Local User Database settings described below. | Yes | No | Yes |

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Users | Controls access to the **Local User Database > Users** node. If you disable this privilege, the administrator will not see the **Local User Database > Users** node or set up a local database on the firewall to store authentication information for remote access users, device administrators, and captive portal users.<br>If you set this privilege to read only, the administrator can view the **Local User Database > Users** information but cannot set up a local database on the firewall to store authentication information. | Yes | Yes | Yes |
| User Groups | Controls access to the **Local User Database > Users** node. If you disable this privilege, the administrator will not see the **Local User Database > Users** node or be able to add user group information to the local database.<br>If you set this privilege to read only, the administrator can view the **Local User Database > Users** information but cannot add user group information to the local database. | Yes | Yes | Yes |
| Authentication Profile | Controls access to the **Authentication Profile** node. If you disable this privilege, the administrator will not see the **Authentication Profile** node or be able to create or edit authentication profiles that specify local database, RADIUS, LDAP, or Kerberos settings that can be assigned to administrator accounts.<br>If you set this privilege to read only, the administrator can view the **Authentication Profile** information but cannot create or edit an authentication profile. | Yes | Yes | Yes |
| Authentication Sequence | Controls access to the **Authentication Sequence** node. If you disable this privilege, the administrator will not see the **Authentication Sequence** node or be able to create or edit an authentication sequence.<br>If you set this privilege to read only, the administrator can view the **Authentication Profile** information but cannot create or edit an authentication sequence. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Access Domain | Controls access to the **Authentication Sequence** node. If you disable this privilege, the administrator will not see the **Authentication Sequence** node or be able to create or edit an authentication sequence.<br>If you set this privilege to read only, the administrator can view the **Authentication Profile** information but cannot create or edit an authentication sequence. | Yes | Yes | Yes |
| Scheduled Log Export | Controls access to the **Scheduled Log Export** node. If you disable this privilege, the administrator will not see the **Scheduled Log Export** node or be able schedule exports of logs and save them to a File Transfer Protocol (FTP) server in CSV format or use Secure Copy (SCP) to securely transfer data between the device and a remote host.<br>If you set this privilege to read only, the administrator can view the **Scheduled Log Export Profile** information but cannot schedule the export of logs. | Yes | No | Yes |
| Software | Controls access to the **Software** node. If you disable this privilege, the administrator will not see the **Software** node or view the latest versions of the PAN-OS software available from Palo Alto Networks, read the release notes for each version, and select a release to download and install.<br>If you set this privilege to read only, the administrator can view the **Software** information but cannot download or install software. | Yes | Yes | Yes |
| GlobalProtect Client | Controls access to the **GlobalProtect Client** node. If you disable this privilege, the administrator will not see the **GlobalProtect Client** node or view available GlobalProtect releases, download the code or activate the GlobalProtect agent.<br>If you set this privilege to read only, the administrator can view the available **GlobalProtect Client** releases but cannot download or install the agent software. | Yes | Yes | Yes |

| Access Level | | Enable | Read Only | Disable |
|---|---|---|---|---|
| Dynamic Updates | Controls access to the **Dynamic Updates** node. If you disable this privilege, the administrator will not see the **Dynamic Updates** node or be able to view the latest updates, read the release notes for each update, or select an update to upload and install.<br>If you set this privilege to read only, the administrator can view the available **Dynamic Updates** releases, read the release notes but cannot upload or install the software. | Yes | Yes | Yes |
| Licenses | Controls access to the **Licenses** node. If you disable this privilege, the administrator will not see the **Licenses** node or be able to view the licenses installed or activate licenses.<br>If you set this privilege to read only, the administrator can view the installed **Licenses**, but cannot perform license management functions. | Yes | Yes | Yes |
| Support | Controls access to the **Support** node. If you disable this privilege, the administrator will not see the **Support** node or be able to access product and security alerts from Palo Alto Networks or generate tech support or stats dump files.<br>If you set this privilege to read only, the administrator can view the **Support** node and access product and security alerts but cannot generate tech support or stats dump files. | Yes | Yes | Yes |
| Master Key and Diagnostics | Controls access to the **Master Key and Diagnostics** node. If you disable this privilege, the administrator will not see the **Master Key and Diagnostics** node or be able to specify a master key to encrypt private keys on the firewall.<br>If you set this privilege to read only, the administrator can view the **Master Key and Diagnostics** node and view information about master keys that have been specified but cannot add or edit a new master key configuration. | Yes | Yes | Yes |

## Define User Privacy Settings in the Admin Role Profile

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Privacy | Sets the default state to enable or disable for all of the privacy settings described below. | Yes | N/A | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Show Full IP addresses | When set to disable, full IP addresses obtained by traffic running through the Palo Alto firewall are not shown in logs or reports. In place of the IP addresses that are normally displayed, the relevant subnet is displayed.<br> Scheduled reports that are displayed in the interface through **Monitor> Reports** and reports that are sent via scheduled emails will still display full IP addresses. Because of this exception, we recommend that the following settings within the **Monitor** tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler. | Yes | N/A | Yes |
| Show User Names in Logs and Reports | When set to disable, user names obtained by traffic running through the Palo Alto Networks firewall are not shown in logs or reports. Columns where the user names would normally be displayed are empty.<br> Scheduled reports that are displayed in the interface through **Monitor> Reports** or reports that are sent via the email scheduler will still display user names. Because of this exception, we recommend that the following settings within the Monitor tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler. | Yes | N/A | Yes |
| View Pcap Files | When set to disable, packet capture files that are normally available within the Traffic, Threat and Data Filtering logs are not displayed. | Yes | N/A | Yes |

## Restrict Admin Access to Commit Functions

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Commit | When set to disable, an admin cannot commit any changes to a configuration. | Yes | N/A | Yes |

### Provide Granular Access to Global Settings

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Global | Sets the default state to enable or disable for all of the global settings described below. In effect, this setting is only for System Alarms at this time. | Yes | N/A | Yes |
| System Alarms | When set to disable, an admin cannot view or acknowledge alarms that are generated. | Yes | N/A | Yes |

## Panorama Web Interface Access

On Panorama the Admin Roles allow you to define access to the options on Panorama and the ability to only allow access to Device Group and Template (**Policies**, **Objects**, **Network**, **Device** tabs).

The admin roles that you can create are: **Panorama** and **Device Group and Template**. The **Device Group and Template** admin role does not provide CLI access privileges.

If an administrator is given superuser privileges on the CLI, the administrator has complete access to all features regardless of the privileges given from the web interface.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Dashboard | Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets. | Yes | No | Yes |
| ACC | Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Yes | No | Yes |
| Monitor | Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the admin can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Policies | Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the admin can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab. | Yes | No | Yes |
| Objects | Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab. | Yes | No | Yes |
| Network | Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the admin can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab. | Yes | No | Yes |
| Device | Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any device-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab.<br> You cannot enable access to the **Admin Roles** or **Administrators** nodes for a role-based administrator even if you enable full access to the **Device** tab. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

# Certificate Management

The following topics describe the different keys and certificates that Palo Alto Networks devices use, and how to obtain and manage them:

- Keys and Certificates
- Certificate Revocation
- Certificate Deployment
- Set Up Verification for Certificate Revocation Status
- Configure the Master Key
- Obtain Certificates
- Configure a Certificate Profile
- Revoke and Renew Certificates
- Secure Keys with a Hardware Security Module

© Palo Alto Networks, Inc.

# Keys and Certificates

To ensure trust between parties in a secure communication session, Palo Alto Networks devices use digital certificates. Each certificate contains a cryptographic key to encrypt plaintext or decrypt cyphertext. Each certificate also includes a digital signature to authenticate the identity of the issuer. The issuer must be in the list of trusted certificate authorities (CAs) of the authenticating party. Optionally, the authenticating party verifies the issuer did not revoke the certificate (see Certificate Revocation).

Palo Alto Networks devices use certificates in the following applications:

- User authentication for Captive Portal, GlobalProtect, Mobile Security Manager, and firewall/Panorama web interface access.
- Device authentication for GlobalProtect VPN (remote user-to-site or large scale).
- Device authentication for IPSec site-to-site VPN with Internet Key Exchange (IKE).
- Decrypting inbound and outbound SSL traffic. A firewall decrypts the traffic to apply security policies and rules, then re-encrypts it before forwarding the traffic to the final destination. For outbound traffic, the firewall acts as a forward proxy server, establishing an SSL/TLS connection to the destination server. To secure a connection between itself and the client, the firewall uses a *signing certificate* to automatically generate a copy of the destination server certificate.

The following table describes the keys and certificates that Palo Alto Networks devices use. As a best practice, use different keys/certificates for each usage.

**Table: Palo Alto Networks Device Keys/Certificates**

| Key/Certificate Usage | Description |
|---|---|
| Administrative Access | Secure access to device administration interfaces (HTTPS access to the web interface) requires a server certificate for the MGT interface (or a designated interface on the dataplane if the device does not use MGT) and, optionally, a certificate to authenticate the administrator. |
| Captive Portal | In deployments where Captive Portal identifies users who access HTTPS resources, designate a server certificate for the Captive Portal interface. If you configure Captive Portal to use certificates (instead of, or in addition to, username/password credentials) for user identification, designate a user certificate also. For more information on Captive Portal, see Map IP Addresses to User Names Using Captive Portal. |
| Forward Trust | For outbound SSL/TLS traffic, if a firewall acting as a forward proxy trusts the CA that signed the certificate of the destination server, the firewall uses the forward trust CA certificate to generate a copy of the destination server certificate to present to the client. The firewall uses the same decryption key for all forward trust certificates. For added security, store the key on a hardware security module (for details, see Secure Keys with a Hardware Security Module). |
| Forward Untrust | For outbound SSL/TLS traffic, if a firewall acting as a forward proxy does not trust the CA that signed the certificate of the destination server, the firewall uses the forward untrust CA certificate to generate a copy of the destination server certificate to present to the client. |

 © Palo Alto Networks, Inc.

| Key/Certificate Usage | Description |
|---|---|
| SSL Inbound Inspection | The keys that decrypt inbound SSL/TLS traffic for inspection and policy enforcement. For this application, import onto the firewall a private key for each server that is subject to SSL/TLS inbound inspection. See Configure SSL Inbound Inspection. |
| SSL Exclude Certificate | Certificates for servers to exclude from SSL/TLS decryption. For example, if you enable SSL decryption but your network includes servers for which the firewall should not decrypt traffic (for example, web services for your HR systems), import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Decryption Exceptions. |
| GlobalProtect | All interaction among GlobalProtect components occurs over SSL/TLS connections. Therefore, as part of the GlobalProtect deployment, deploy server certificates for all GlobalProtect portals, gateways, and Mobile Security Managers. Optionally, deploy certificates for authenticating users also.<br>Note that the GlobalProtect Large Scale VPN (LSVPN) feature requires a CA signing certificate. |
| Site-to-Site VPNs (IKE) | In a site-to-site IPSec VPN deployment, peer devices use Internet Key Exchange (IKE) gateways to establish a secure channel. IKE gateways use certificates or preshared keys to authenticate the peers to each other. You configure and assign the certificates or keys when defining an IKE gateway on a firewall. See Site-to-Site VPN Overview. |
| Master Key | The firewall uses a master key to encrypt all private keys and passwords. If your network requires a secure location for storing private keys, you can use an encryption (wrapping) key stored on a hardware security module (HSM) to encrypt the master key. For details, see Encrypt a Master Key Using an HSM. |
| Secure Syslog | The certificate to enable secure connections between the firewall and a syslog server. See Configure the Firewall to Authenticate to the Syslog Server. |
| Trusted Root CA | The designation for a root certificate issued by a CA that the firewall trusts. The firewall can use a self-signed root CA certificate to automatically issue certificates for other applications (for example, SSL Forward Proxy).<br>Also, if a firewall must establish secure connections with other firewalls, the root CA that issues their certificates must be in the list of trusted root CAs on the firewall. |

© Palo Alto Networks, Inc.

# Certificate Revocation

Palo Alto Networks devices use digital certificates to ensure trust between parties in a secure communication session. Configuring a device to check the revocation status of certificates provides additional security. A party that presents a revoked certificate is not trustworthy. When a certificate is part of a chain, the device checks the status of every certificate in the chain except the root CA certificate, for which the device cannot verify revocation status.

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority that issued the certificate must revoke it.

Palo Alto Networks devices support the following methods for verifying certificate revocation status. If you configure both, the devices first try the OCSP method; if the OCSP server is unavailable, the devices use the CRL method.

- Certificate Revocation List (CRL)
- Open Certificate Status Protocol (OCSP)

In PAN-OS, certificate revocation status verification is an optional feature. It is a best practice to enable it for certificate profiles, which define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPsec VPN, and firewall/Panorama web interface access.

## Certificate Revocation List (CRL)

Each certificate authority (CA) periodically issues a certificate revocation list (CRL) to a public repository. The CRL identifies revoked certificates by serial number. After the CA revokes a certificate, the next CRL update will include the serial number of that certificate.

The Palo Alto Networks firewall downloads and caches the last-issued CRL for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the CRL for the issuing CA. Also, the cache only stores a CRL until it expires.

To use CRLs for verifying the revocation status of certificates when the firewall functions as an SSL forward proxy, Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

To use CRLs for verifying the revocation status of certificates that authenticate users and devices, configure a certificate profile and assign it to the interfaces that are specific to the application: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, or firewall/Panorama web interface access. For details, see Configure Revocation Status Verification of Certificates Used for User/Device Authentication.

## Open Certificate Status Protocol (OCSP)

When establishing an SSL/TLS session, clients can use Online Certificate Status Protocol (OCSP) to check the revocation status of the authentication certificate. The authenticating client sends a request containing the serial number of the certificate to the OCSP responder (server). The responder searches the database of the certificate

 © Palo Alto Networks, Inc.

authority (CA) that issued the certificate and returns a response containing the status (*good*, *revoked* or *unknown*) to the client. The advantage of the OCSP method is that it can verify status in real-time, instead of depending on the issue frequency (hourly, daily, or weekly) of CRLs.

The Palo Alto Networks firewall downloads and caches OCSP status information for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the OCSP information for the issuing CA. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall as an OCSP responder (see Configure an OCSP Responder).

To use OCSP for verifying the revocation status of certificates when the firewall functions as an SSL forward proxy, perform the steps under Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

The following applications use certificates to authenticate users and/or devices: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, and firewall/Panorama web interface access. To use OCSP for verifying the revocation status of the certificates:

- Configure an OCSP responder.
- Enable the HTTP OCSP service on the firewall.
- Create or obtain a certificate for each application.
- Configure a certificate profile for each application.
- Assign the certificate profile to the relevant application.

To cover situations where the OCSP responder is unavailable, configure CRL as a fall-back method. For details, see Configure Revocation Status Verification of Certificates Used for User/Device Authentication.

# Certificate Deployment

The basic approaches to deploy certificates for Palo Alto Networks devices are:

- **Obtain certificates from a trusted third-party CA**—The benefit of obtaining a certificate from a trusted third-party certificate authority (CA) such as VeriSign or GoDaddy is that end clients will already trust the certificate because common browsers include root CA certificates from well-known CAs in their trusted root certificate stores. Therefore, for applications that require end clients to establish secure connections with a Palo Alto Network device, purchase a certificate from a CA that the end clients trust to avoid having to pre-deploy root CA certificates to the end clients. (Some such applications are a GlobalProtect portal or GlobalProtect Mobile Security Manager.) However, note that most third-party CAs cannot issue signing certificates. Therefore, this type of certificate is not appropriate for applications (for example, SSL/TLS decryption and large-scale VPN) that require the firewall to issue certificates.
- **Obtain certificates from an enterprise CA**—Enterprises that have their own internal CA can use it to issue certificates for firewall applications and import them onto the firewall. The benefit is that end clients probably already trust the enterprise CA. You can either generate the needed certificates and import them onto the firewall, or generate a certificate signing request (CSR) on the firewall and send it to the enterprise CA for signing. The benefit of this method is that the private key does not leave the firewall. An enterprise CA can also issue a signing certificate, which the firewall uses to automatically generate certificates (for example, for GlobalProtect large-scale VPN or sites requiring SSL/TLS decryption).
- **Generate self-signed certificates**—You can generate a self-signed root CA certificate on the firewall and use it to automatically issue certificates for other firewall applications. Note that if you use this method to generate certificates for an application that requires an end client to trust the certificate, end users will see a certificate error because the root CA certificate is not in their trusted root certificate store. To prevent this, deploy the self-signed root CA certificate to all end user systems. You can deploy the certificates manually or use a centralized deployment method such as an Active Directory Group Policy Object (GPO).

 © Palo Alto Networks, Inc.

# Set Up Verification for Certificate Revocation Status

To verify the revocation status of certificates, the firewall uses Open Certificate Status Protocol (OCSP) and/or certificate revocation lists (CRLs). For details on these methods, see Certificate Revocation If you configure both methods, the firewall first tries OCSP and only falls back to the CRL method if the OCSP responder is unavailable. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall to function as the OCSP responder.

The following topics describe how to configure the firewall to verify certificate revocation status:

- Configure an OCSP Responder
- Configure Revocation Status Verification of Certificates Used for User/Device Authentication
- Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

## Configure an OCSP Responder

To use Open Certificate Status Protocol (OCSP) for verifying the revocation status of certificates, you must configure the firewall to access an OCSP responder (server). The entity that manages the OCSP responder can be a third-party certificate authority (CA) or, if your enterprise has its own public key infrastructure (PKI), the firewall itself. For details on OCSP, see Certificate Revocation

| **Configure an OCSP Responder** | |
|---|---|
| Step 1 Define an OCSP responder. | 1. In a firewall, select **Device > Certificate Management > OCSP Responder** and click **Add**.<br>In Panorama, select **Device > Certificate Management > OCSP Responder**, select a **Template** and click **Add**.<br>2. Enter a **Name** to identify the responder (up to 31 characters). The name is case-sensitive. It must be unique and use only letters, numbers, spaces, hyphens, and underscores.<br>3. If the firewall supports multiple virtual systems, the dialog displays a **Location** drop-down. Select the virtual system where the responder will be available or select **Shared** to enable availability on all the virtual systems.<br>4. In the **Host Name** field, enter the host name (recommended) or IP address of the OCSP responder. From this value, PAN-OS automatically derives a URL and adds it to the certificate being verified.<br>If you configure the firewall itself as an OCSP responder, the host name must resolve to an IP address in the interface that the firewall uses for OCSP services (specified in Step 3).<br>5. Click **OK**. |

| Configure an OCSP Responder | | |
|---|---|---|
| Step 2 | Enable OCSP communication on the firewall. | 1. In a firewall, select **Device > Setup > Management**.<br>In Panorama, select **Device > Setup > Management** and select a **Template**.<br>2. In the Management Interface Settings section, click the Edit icon, select the **HTTP OCSP** check box, then click **OK**. |
| Step 3 | Optionally, to configure the firewall itself as an OCSP responder, add an Interface Management Profile to the interface used for OCSP services. | 1. Select **Network > Network Profiles > Interface Mgmt**.<br>2. Click **Add** to create a new profile or click the name of an existing profile.<br>3. Select the **HTTP OCSP** check box and click **OK**.<br>4. Select **Network > Interfaces** and click the name of the interface that the firewall will use for OCSP services. The OCSP **Host Name** specified in Step 1 must resolve to an IP address in this interface.<br>5. Select **Advanced > Other info** and select the Interface Management Profile you configured.<br>6. Click **OK** and **Commit**. |

## Configure Revocation Status Verification of Certificates Used for User/Device Authentication

The firewall uses certificates to authenticate users and devices for such applications as Captive Portal, GlobalProtect, site-to-site IPSec VPN, and firewall/Panorama web interface access. To improve security, it is a best practice to configure the firewall to verify the revocation status of certificates that it uses for device/user authentication.

| Configure Revocation Status Verification of Certificates Used for User/Device Authentication | | |
|---|---|---|
| Step 1 | Configure a Certificate Profile for each application. | Assign one or more root CA certificates to the profile and select how the firewall verifies certificate revocation status. The common name (FQDN or IP address) of a certificate must match an interface to which you apply the profile in Step 2.<br>For details on the certificates that various applications use, see Keys and Certificates |
| Step 2 | Assign the certificate profiles to the relevant applications. | The steps to assign a certificate profile depend on the application that requires it. |

 © Palo Alto Networks, Inc.

# Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

The firewall decrypts outbound SSL/TLS traffic to apply security policies and rules, then re-encrypts the traffic before forwarding it. During this process, the firewall acts as a forward proxy server, maintaining separate connections with the client and destination server. For the client connection, the firewall uses a CA certificate to automatically generate a decryption certificate that is a copy of the destination server certificate. You can configure the firewall to verify the revocation status of destination server certificates as follows.

Enabling revocation status verification for SSL/TLS decryption certificates will add time to the process of establishing the session. The first attempt to access a site might fail if the verification does not finish before the session times out. For these reasons, verification is disabled by default.

| **Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption** | | |
|---|---|---|
| Step 1 | Access the Decryption Certificate Revocation Settings page. | In a firewall, select **Device > Setup > Session** and, in the Session Features section, select **Decryption Certificate Revocation Settings**.<br>In Panorama, select **Device > Setup > Session**, select a **Template** and, in the Session Features section, select **Decryption Certificate Revocation Settings**. |
| Step 2 | Define the service-specific timeout intervals for revocation status requests. | Perform one or both of the following steps, depending on whether the firewall will use Open Certificate Status Protocol (OCSP) or the certificate revocation list (CRL) method to verify the revocation status of certificates. If the firewall will use both, it first tries OCSP; if the OCSP responder is unavailable, the firewall then tries the CRL method.<br>1. In the CRL section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the CRL service.<br>2. In the OCSP section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the OCSP responder.<br>Depending on the **Certificate Status Timeout** value you specify in Step 3, the firewall might register a timeout before either or both of the **Receive Timeout** intervals pass. |

| Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption | | |
|---|---|---|
| Step 3 | Define the total timeout interval for revocation status requests. | Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies the session blocking logic you optionally define in Step 4. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:<br>• If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.<br>• If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.<br>• If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value. |
| Step 4 | Define the blocking behavior for *unknown* certificate status or a revocation status request timeout. | If you want the firewall to block SSL/TLS sessions when the OCSP or CRL service returns a certificate revocation status of *unknown*, select the **Block Session With Unknown Certificate Status** check box. Otherwise, the firewall proceeds with the session.<br>If you want the firewall to block SSL/TLS sessions after it registers a request timeout, select the **Block Session On Certificate Status Check Timeout** check box. Otherwise, the firewall proceeds with the session. |
| Step 5 | Save and apply your entries. | Click **OK** and **Commit**. |


© Palo Alto Networks, Inc.

# Configure the Master Key

Every firewall has a default master key that encrypts private keys and other secrets (such as passwords and shared keys). The private keys authenticate users when they access administrative interfaces on the firewall. As a best practice to safeguard the keys, configure the master key on each firewall to be unique and periodically change it. For added security, use a *wrapping key* stored on a hardware security module (HSM) to encrypt the master key. For details, see Encrypt a Master Key Using an HSM.

In a high availability (HA) configuration, ensure both devices in the pair use the same master key to encrypt private keys and certificates. If the master keys differ, HA configuration synchronization will not work properly.

When you export a firewall configuration, the master key encrypts the passwords of users managed on external servers. For locally managed users, the firewall hashes the passwords but the master key does not encrypt them.

| Configure a Master Key |
|---|
| 1. In a firewall, select **Device > Master Key and Diagnostics** and, in the Master Key section, click the Edit icon.<br>In Panorama, select **Panorama > Master Key and Diagnostics** and, in the Master Key section, click the Edit icon. |
| 2. Enter the **Current Master Key** if one exists. |
| 3. Define a new **New Master Key** and then **Confirm New Master Key**. The key must contain exactly 16 characters. |
| 4. (Optional) To specify the master key **Life Time**, enter the number of **Days** and/or **Hours** after which the key will expire. If you set a life time, create a new master key before the old key expires. |
| 5. (Optional) If you set a key life time, enter a **Time for Reminder** that specifies the number of **Days** and **Hours** preceding master key expiration when the firewall emails you a reminder. |
| 6. (Optional) Select whether to use an **HSM** to encrypt the master key. For details, see Encrypt a Master Key Using an HSM. |
| 7. Click **OK** and **Commit**. |

© Palo Alto Networks, Inc.

# Obtain Certificates


- Create a Self-Signed Root CA Certificate
- Generate a Certificate on the Firewall
- Import a Certificate and Private Key
- Obtain a Certificate from an External CA


## Create a Self-Signed Root CA Certificate

A self-signed root certificate authority (CA) certificate is the top-most certificate in a certificate chain. A firewall can use this certificate to automatically issue certificates for other uses. For example, the firewall issues certificates for SSL/TLS decryption and for satellite devices in a GlobalProtect large-scale VPN.

When establishing a secure connection with the firewall, the remote client must trust the root CA that issued the certificate. Otherwise, the client browser will display a warning that the certificate is invalid and might (depending on security settings) block the connection. To prevent this, after generating the self-signed root CA certificate, import it into the client systems.

| **Generate a Self-signed Root CA Certificate** |
|---|
| 1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.<br>In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**. |
| 2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in Step 6. |
| 3. Click **Generate**. |
| 4. Enter a **Certificate Name**, such as *GlobalProtect_CA*. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores. |
| 5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate. |
| 6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box only appears if the device supports multiple virtual systems. |
| 7. Leave the **Signed By** field blank to designate the certificate as self-signed. |
| 8. Select the **Certificate Authority** check box. |
| 9. Do *not* select an **OCSP Responder**. Certificate revocation status verification does not apply to root CA certificates. |
| 10. Click **Generate** and **Commit**. |

 © Palo Alto Networks, Inc.

## Generate a Certificate on the Firewall

The firewall uses certificates to authenticate clients, servers, users and devices in several applications, including SSL/TLS decryption, Captive Portal, GlobalProtect, site-to-site IPSec VPN, and firewall/Panorama web interface access. Generate certificates for each usage. For details on application-specific certificates, see Keys and Certificates

To generate a certificate, you must first create or import a root CA certificate to sign it. For details, see Create a Self-Signed Root CA Certificate or Import a Certificate and Private Key.

To use Open Certificate Status Protocol (OCSP) for verifying certificate revocation status, Configure an OCSP Responder *before* generating the certificate. For details on status verification, see Certificate Revocation

**Generate a Certificate on the Firewall**

1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.
   In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.
2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in Step 6.
3. Click **Generate**.
4. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.
5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.
6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box only appears if the device supports multiple virtual systems.
7. In the **Signed By** field, select the root CA certificate that will issue the certificate.
8. If applicable, select an **OCSP Responder**.
9. (Optional) Define the **Cryptographic Settings** as necessary to create a certificate that will work with the devices that must authenticate to it. The default and recommended key size (**Number of Bits**) is 2048 bits. The default and recommended encryption algorithm (**Digest**) is SHA256.
10. (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.

    If you add a **Host Name** (DNS name) attribute, it is a best practice for it to match the **Common Name**. The host name populates the Subject Alternative Name field of the certificate.
11. Click **Generate** and, in the Device Certificates tab, click the certificate **Name**.

    Regardless of the time zone on the firewall, it always displays the corresponding Greenwich Mean Time (GMT) for certificate validity and expiration dates/times.
12. Select the check boxes that correspond to the intended use of the certificate on the firewall. For example, if the firewall will use this certificate to authenticate user access to its web interface, select the **Certificate for Secure Web GUI** check box.
13. Click **OK** and **Commit**.

## Import a Certificate and Private Key

If your enterprise has its own public key infrastructure (PKI), you can import a certificate and private key into the firewall from your enterprise certificate authority (CA). Enterprise CA certificates (unlike most certificates purchased from a trusted, third-party CA) can automatically issue CA certificates for applications such as SSL/TLS decryption or large-scale VPN.

Instead of importing a self-signed root CA certificate into all the client systems, it is a best practice to import a certificate from the enterprise CA because the clients will already have a trust relationship with the enterprise CA, which simplifies the deployment.

If the certificate you will import is part of a certificate chain, it is a best practice to import the entire chain.

**Import a Certificate and Private Key**

1. From the enterprise CA, export the certificate and private key that the firewall will use for authentication.

   When exporting a private key, you must enter a passphrase to encrypt the key for transport. Ensure the management system can access the certificate and key files. When importing the key onto the firewall, you must enter the same passphrase to decrypt it.

2. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.

   In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.

3. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in Step 6.

4. Click **Import**.

5. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.

6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box only appears if the device supports multiple virtual systems.

7. Enter the path and name of the **Certificate File** received from the CA, or **Browse** to find the file.

8. Select a **File Format**:
   - **Encrypted Private Key and Certificate (PKCS12)**—This is the default and most common format, in which the key and certificate are in a single container (**Certificate File**). If a hardware security module (HSM) will store the private key for this certificate, select the **Private key resides on Hardware Security Module** check box.
   - **Base64 Encoded Certificate (PEM)**—You must import the key separately from the certificate. If a hardware security module (HSM) stores the private key for this certificate, select the **Private key resides on Hardware Security Module** check box and skip Step 9. Otherwise, select the **Import Private Key** checkbox, enter the **Key File** or **Browse** to it, then perform Step 9.

9. Enter and re-enter (confirm) the **Passphrase** used to encrypt the private key.

10. Click **OK**. The Device Certificates tab displays the imported certificate.

 © Palo Alto Networks, Inc.

## Obtain a Certificate from an External CA

The advantage of obtaining a certificate from an external certificate authority (CA) is that the private key does not leave the firewall. To obtain a certificate from an external CA, generate a certificate signing request (CSR) and submit it to the CA. After the CA issues a certificate with the specified attributes, import it onto the firewall. The CA can be a well-known, public CA or an enterprise CA.

To use Open Certificate Status Protocol (OCSP) for verifying the revocation status of the certificate, Configure an OCSP Responder *before* generating the CSR.

| **Obtain a Certificate from an External CA** | | |
|---|---|---|
| Step 1 | Request the certificate from an external CA. | 1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.<br>In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.<br>2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in sub-step 6.<br>3. Click **Generate**.<br>4. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.<br>5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.<br>6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box only appears if the device supports multiple virtual systems.<br>7. In the **Signed By** field, select **External Authority (CSR)**.<br>8. If applicable, select an **OCSP Responder**.<br>9. (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.<br>If you add a **Host Name** attribute, it is a best practice for it to match the **Common Name** (this is mandatory for GlobalProtect). The host name populates the Subject Alternative Name field of the certificate.<br>10. Click **Generate**. The Device Certificates tab displays the CSR with a Status of *pending*. |
| Step 2 | Submit the CSR to the CA. | 1. Select the CSR and click **Export** to save the .csr file to a local computer.<br>2. Upload the .csr file to the CA. |

| Obtain a Certificate from an External CA | |
|---|---|
| Step 3 Import the certificate. | 1. After the CA sends a signed certificate in response to the CSR, return to the Device Certificates tab and click **Import**.<br>2. Enter the **Certificate Name** used to generate the CSR in Step 1-4.<br>3. Enter the path and name of the PEM **Certificate File** that the CA sent, or **Browse** to it.<br>4. Click **OK**. The Device Certificates tab displays the certificate with a Status of *valid*. |
| Step 4 Configure the certificate. | 1. Click the certificate **Name**.<br>2. Select the check boxes that correspond to the intended use of the certificate on the firewall. For example, if the firewall will use this certificate to authenticate administrators who access the web interface, select the **Certificate for Secure Web GUI** check box.<br>3. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# Configure a Certificate Profile

Certificate profiles define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPSec VPN, Mobile Security Manager, and firewall/Panorama web interface access. The profiles specify which certificates to use, how to verify certificate revocation status, and how that status constrains access. Configure a certificate profile for each application.

It is a best practice to enable Open Certificate Status Protocol (OCSP) and/or Certificate Revocation List (CRL) status verification for certificate profiles. For details on these methods, see Certificate Revocation

| Configure a Certificate Profile | | |
|---|---|---|
| Step 1 | Obtain the certificate authority (CA) certificates you will assign. | Perform one of the following steps to obtain the CA certificates you will assign to the profile. You must assign at least one.<br>• Create a Self-Signed Root CA Certificate.<br>• Export a certificate from your enterprise CA and then import it onto the firewall (see Step 3). |
| Step 2 | Identify the certificate profile. | 1. In a firewall, select **Device > Certificate Management > Certificates Profile** and click **Add**.<br>In Panorama, select **Device > Certificate Management > Certificates Profile**, select a **Template** and click **Add**.<br>2. Enter a **Name** to identify the profile. The name is case-sensitive. It must be unique and use only letters, numbers, spaces, hyphens, and underscores. It can have up to 31 characters.<br>3. If the firewall supports multiple virtual systems, the dialog displays a **Location** drop-down. Select the virtual system where the profile will be available or select **Shared** to enable availability on all the virtual systems. |

| Configure a Certificate Profile | |
|---|---|
| Step 3 Assign one or more certificates. | Perform the following steps for each certificate:<br>1. In the CA Certificates table, click **Add**.<br>2. Select a **CA Certificate** from Step 1, or click **Import** and perform the following sub-steps:<br>a. Enter a **Certificate Name**.<br>b. Enter the path and name of the **Certificate File** you exported from your enterprise CA, or **Browse** to find the file.<br>c. Click **OK**.<br>3. Optionally, if the firewall uses OCSP to verify certificate revocation status, configure the following fields to override the default behavior. For most deployments, these fields do not apply.<br>• By default, the firewall uses the OCSP responder URL that you set in the procedure Configure an OCSP Responder. To override that setting, enter a **Default OCSP URL** (starting with *http://* or *https://*).<br>• By default, the firewall uses the certificate selected in the **CA Certificate** field to validate OCSP responses. To use a different certificate for validation, select it in the **OCSP Verify CA Certificate** field.<br>4. Click **OK**. The CA Certificates table displays the assigned certificate. |


© Palo Alto Networks, Inc.

| Configure a Certificate Profile | | |
|---|---|---|
| Step 4 | Define the methods for verifying certificate revocation status and the associated blocking behavior. | 1. Select **Use CRL** and/or **Use OCSP**. If you select both, the firewall first tries OCSP and only falls back to the CRL method if the OCSP responder is unavailable.<br>2. Depending on the verification method, enter the **CRL Receive Timeout** and/or **OCSP Receive Timeout**. These are the intervals (1-60 seconds) after which the firewall stops waiting for a response from the CRL/OCSP service.<br>3. Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies any session blocking logic you define. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:<br>• If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.<br>• If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.<br>• If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value.<br>4. If you want the firewall to block sessions when the OCSP or CRL service returns a certificate revocation status of *unknown*, select the **Block session if certificate status is unknown** check box. Otherwise, the firewall proceeds with the session.<br>5. If you want the firewall to block sessions after it registers an OCSP or CRL request timeout, select the **Block session if certificate status cannot be retrieved within timeout** check box. Otherwise, the firewall proceeds with the session. |
| Step 5 | Save and apply your entries. | Click **OK** and **Commit**. |

# Revoke and Renew Certificates

▲ Revoke a Certificate
▲ Renew a Certificate

## Revoke a Certificate

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority (CA) that issued the certificate must revoke it. The following task describes how to revoke a certificate for which the firewall is the CA.

| **Revoke a Certificate** |
|---|
| 1. Select **Device > Certificate Management > Certificates > Device Certificates**. |
| 2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select the virtual system to which the certificate belongs. |
| 3. Select the certificate to revoke. |
| 4. Click **Revoke**. PAN-OS immediately sets the status of the certificate to revoked and adds the serial number to the Open Certificate Status Protocol (OCSP) responder cache or certificate revocation list (CRL). You need not perform a commit. |

## Renew a Certificate

If a certificate expires, or soon will, you can reset the validity period. If an external certificate authority (CA) signed the certificate and the firewall uses the Open Certificate Status Protocol (OCSP) to verify certificate revocation status, the firewall uses the OCSP responder information to update the certificate status (see Configure an OCSP Responder). If the firewall is the CA that issued the certificate, the firewall replaces it with a new certificate that has a different serial number but the same attributes as the old certificate.

| **Renew a Certificate** |
|---|
| 1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.<br>In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**. |
| 2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select the virtual system to which the certificate belongs. |
| 3. Select a certificate to renew and click **Renew**. |
| 4. Enter a **New Expiration Interval** (in days). |
| 5. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# Secure Keys with a Hardware Security Module

A hardware security module (HSM) is a physical device that manages digital keys. An HSM provides secure storage and generation of digital keys. It provides both logical and physical protection of these materials from non-authorized use and potential adversaries.

HSM clients integrated with Palo Alto Networks devices enable enhanced security for the private keys used in SSL/TLS decryption (both SSL forward proxy and SSL inbound inspection). In addition, you can use the HSM to encrypt device master keys.

The following topics describe how to integrate an HSM with your Palo Alto Networks devices:

- Set up Connectivity with an HSM
- Encrypt a Master Key Using an HSM
- Store Private Keys on an HSM
- Manage the HSM Deployment

# Set up Connectivity with an HSM

HSM clients are integrated with PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7050, and VM-Series firewalls and on Panorama (virtual appliance and M-100 appliance) for use with the following HSMs:

- SafeNet Luna SA 5.2.1 or later
- Thales Nshield Connect 11.62 or later

The HSM server version must be compatible with these client versions. Refer to the HSM vendor documentation for the client-server version compatibility matrix.

The following topics describe how to set up connectivity between the firewall/Panorama and one of the supported HSMs:

▲ Set Up Connectivity with a SafeNet Luna SA HSM

▲ Set Up Connectivity with a Thales Nshield Connect HSM

## Set Up Connectivity with a SafeNet Luna SA HSM

To set up connectivity between the Palo Alto Networks device and a SafeNet Luna SA HSM, you must specify the address of the HSM server and the password for connecting to it in the firewall configuration. In addition, you must register the firewall with the HSM server. Prior to beginning the configuration, make sure you have created a partition for the Palo Alto Networks devices on the HSM server.

HSM configuration is not synced between high availability firewall peers. Consequently, you must configure the HSM module separately on each of the peers.

In Active-Passive HA deployments, you must manually perform one failover to configure and authenticate each HA peer individually to the HSM. After this manual failover has been performed, user interaction is not required for the failover function.

| **Set up a Connectivity with a SafeNet Luna SA HSM** | |
|---|---|
| Step 1 Configure the firewall to communicate with the SafeNet Luna SA HSM. | 1. Log in to the firewall web interface and select **Device > Setup > HSM**.<br>2. Edit the Hardware Security Module Provider section and select **Safenet Luna SA** as the **Provider Configured**.<br>3. Click **Add** and enter a **Module Name**. This can be any ASCII string up to 31 characters in length.<br>4. Enter the IPv4 address of the HSM module as the **Server Address**.<br>If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices.<br>5. (Optional) If configuring a high availability HSM configuration, select the **High Availability** check box and add the following: a value for **Auto Recovery Retry** and a **High Availability Group Name**.<br>If two HSM servers are configured, you should configure high availability. Otherwise the second HSM server is not used.<br>6. Click **OK** and **Commit.** |

 © Palo Alto Networks, Inc.

| Set up a Connectivity with a SafeNet Luna SA HSM (Continued) | | |
|---|---|---|
| Step 2 | (Optional) Configure a service route to enable the firewall to connect to the HSM.<br><br>By default, the firewall uses the Management Interface to communicate with the HSM. To use a different interface, you must configure a service route. | 1. Select **Device > Setup > Services.**<br>2. Select **Service Route Configuration** from the Services Features area.<br>3. Select **Customize** from the Service Route Configuration area.<br>4. Select the **IPv4** tab.<br>5. Select **HSM** from the **Service** column.<br>6. Select an interface to use for HSM from the **Source Interface** drop-down.<br><br>If you select a dataplane connected port for HSM, issuing the `clear session all` CLI command, will clear all existing HSM sessions causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.<br><br>7. Click **OK** and **Commit.** |
| Step 3 | Configure the firewall to authenticate to the HSM. | 1. Select **Device > Setup > HSM**.<br>2. Select **Setup Hardware Security Module** in the Hardware Security Operations area.<br>3. Select the HSM **Server Name** from the drop-down.<br>4. Enter the **Administrator Password** to authenticate the firewall to the HSM.<br>5. Click **OK.**<br><br>The firewall attempts to perform an authentication with the HSM and displays a status message.<br><br>6. Click **OK.** |
| Step 4 | Register the firewall (the HSM client) with the HSM and assign it to a partition on the HSM.<br><br>If the HSM already has a firewall with the same `<cl-name>` registered, you must remove the duplicate registration using the following command before registration will succeed:<br><br>`client delete -client <cl-name>`<br><br>where `<cl-name>` is the name of the client (firewall) registration you want to delete. | 1. Log in to the HSM from a remote system.<br>2. Register the firewall using the following command:<br><br>`client register -c <cl-name> -ip <fw-ip-addr>`<br><br>where `<cl-name>` is a name that you assign to the firewall for use on the HSM and `<fw-ip-addr>` is the IP address of the firewall that is being configured as an HSM client.<br><br>3. Assign a partition to the firewall using the following command:<br><br>`client assignpartition -c <cl-name> -p <partition-name>`<br><br>where `<cl-name>` is the name assigned to the firewall in the `client register` command and `<partition-name>` is the name of a previously configured partition that you want to assign to the firewall. |

| Set up a Connectivity with a SafeNet Luna SA HSM (Continued) | | |
|---|---|---|
| Step 5 | Configure the firewall to connect to the HSM partition. | 1. Select **Device > Setup > HSM**.<br>2. Click the **Refresh** icon.<br>3. Select the **Setup HSM Partition** in the Hardware Security Operations area.<br>4. Enter the **Partition Password** to authenticate the firewall to the partition on the HSM.<br>5. Click **OK**. |
| Step 6 | (Optional) Configure an additional HSM for high availability (HA). | 1. Follow Step 1 through Step 5 to add an additional HSM for high availability (HA).<br>This process adds a new HSM to the existing HA group.<br>2. If you remove an HSM from your configuration, repeat Step 5.<br>This will remove the deleted HSM from the HA group. |
| Step 7 | Verify connectivity with the HSM. | 1. Select **Device > Setup > HSM**.<br>2. Check the **Status** of the HSM connection:<br>Green = HSM is authenticated and connected<br>Red = HSM was not authenticated or network connectivity to the HSM is down.<br>3. View the following columns in Hardware Security Module Status area to determine authentication status:<br>**Serial Number** —The serial number of the HSM partition if the HSM was successfully authenticated.<br>**Partition**—The partition name on the HSM that was assigned on the firewall.<br>**Module State**—The current operating state of the HSM. It always has the value Authenticated if the HSM is displayed in this table. |

## Set Up Connectivity with a Thales Nshield Connect HSM

The following workflow describes how to configure the firewall to communicate with a Thales Nshield Connect HSM. This configuration requires that you set up a remote filesystem (RFS) to use as a *hub* to sync key data for all firewalls in your organization that are using the HSM.

HSM configuration is not synced between high availability firewall peers. Consequently, you must configure the HSM module separately on each of the peers.

If the high availability firewall configuration is in Active-Passive mode, you must manually perform one failover to configure and authenticate each HA peer individually to the HSM. After this manual failover has been performed, user interaction is not required for the failover function.

 © Palo Alto Networks, Inc.

| Set up Connectivity with a Thales Nshield Connect HSM | | |
|---|---|---|
| Step 1 | Configure the Thales Nshield Connect server as the firewall's HSM provider. | 1. From the firewall web interface, select **Device > Setup > HSM** and edit the Hardware Security Module Provider section.<br>2. Select **Thales Nshield Connect** as the **Provider Configured**.<br>3. Click **Add** and enter a **Module Name**. This can be any ASCII string up to 31 characters in length.<br>4. Enter the IPv4 address as the **Server Address** of the HSM module.<br>If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices.<br>5. Enter the IPv4 address of the **Remote Filesystem Address**.<br>6. Click **OK** and **Commit**. |
| Step 2 | (Optional) Configure a service route to enable the firewall to connect to the HSM.<br>By default, the firewall uses the Management Interface to communicate with the HSM. To use a different interface, you must configure a service route. | 1. Select **Device > Setup > Services.**<br>2. Select **Service Route Configuration** from the Services Features area.<br>3. Select **Customize** from the Service Route Configuration area.<br>4. Select the **IPv4** tab.<br>5. Select **HSM** from the **Service** column.<br>6. Select an interface to use for HSM from the **Source Interface** drop-down.<br>If you select a dataplane connected port for HSM, issuing the `clear session all` CLI command, will clear all existing HSM sessions causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.<br>7. Click **OK** and **Commit.** |
| Step 3 | Register the firewall (the HSM client) with the HSM server.<br>This step briefly describes the procedure for using the front panel interface of the Thales Nshield Connect HSM. For more details, consult the Thales documentation. | 1. Log in to the front panel display of the Thales Nshield Connect HSM unit.<br>2. On the unit front panel, use the right-hand navigation button to select **System > System configuration > Client config > New client**.<br>3. Enter the IP address of the firewall.<br>4. Select **System > System configuration > Client config > Remote file system** and enter the IP address of the client computer where you set up the remote file system. |

| Set up Connectivity with a Thales Nshield Connect HSM (Continued) | | |
|---|---|---|
| Step 4 | Set up the remote filesystem to accept connections from the firewall. | 1. Log in to the remote filesystem (RFS) from a Linux client.<br>2. Obtain the electronic serial number (ESN) and the hash of the $K_{NETI}$ key. The $K_{NETI}$ key authenticates the module to clients:<br>`anonkneti <ip-address>`<br>where `<ip-address>` is the IP address of the HSM.<br>The following is an example:<br>`anonkneti 192.0.2.1 B1E2-2D4C-E6A2 5a2e5107e70d525615a903f6391ad72b1c03352c`<br>In this example, **`B1E2-2D4C-E6A2`** is the ESM and 5a2e5107e70d525615a903f6391ad72b1c03352c is the hash of the $K_{NETI}$ key.<br>3. Use the following command from a superuser account to perform the remote filesystem setup:<br>`rfs-setup --force <ip-address> <ESN> <hash-Kneti-key>`<br>where `<ip-address>` is the IP address of the HSM,<br>`<ESN>` is the electronic serial number (ESN) and<br>`<hash-Kneti-key>` is the hash of the KNETI key.<br>The following example uses the values obtained in this procedure:<br>`rfs-setup --force <192.0.2.1> <B1E2-2D4C-E6A2> <5a2e5107e70d525615a903f6391ad72b1c03352c>`<br>4. Use the following command to permit client submit on the Remote Filesystem:<br>`rfs-setup --gang-client --write-noauth <FW-IPaddress>`<br>where `<FW-IPaddress>` is the IP address of the firewall. |
| Step 5 | Configure the firewall to authenticate to the HSM. | 1. From the firewall web interface, select **Device > Setup > HSM**.<br>2. Select **Setup Hardware Security Module** in the Hardware Security Operations area.<br>3. Click **OK.**<br>The firewall attempts to perform an authentication with the HSM and displays a status message.<br>4. Click **OK.** |
| Step 6 | Synchronize the firewall with the remote filesystem. | 1. Select the **Device > Setup > HSM**.<br>2. Select **Synchronize with Remote Filesystem** in the Hardware Security Operations section. |

 © Palo Alto Networks, Inc.

| Set up Connectivity with a Thales Nshield Connect HSM (Continued) | | |
|---|---|---|
| Step 7 | Verify that the firewall can connect to the HSM. | 1. Select **Device > Setup > HSM**.<br>2. Check the Status indicator to verify that the firewall is connected to the HSM:<br>Green = HSM is authenticated and connected.<br>Red = HSM was not authenticated or network connectivity to the HSM is down.<br>3. View the following columns in Hardware Security Module Status section to determine authentication status.<br>**Name**: The name of the HSM attempting to be authenticated.<br>**IP address:** The IP address of the HSM that was assigned on the firewall.<br>**Module State**: The current operating state of the HSM: **Authenticated** or **Not Authenticated**. |

# Encrypt a Master Key Using an HSM

A master key is configured on a Palo Alto Networks firewall to encrypt all private keys and passwords. If you have security requirements to store your private keys in a secure location, you can encrypt the master key using an encryption key that is stored on an HSM. The firewall then requests the HSM to decrypt the master key whenever it is required to decrypt a password or private key on the firewall. Typically, the HSM is located in a highly secure location that is separate from the firewall for greater security.

The HSM encrypts the master key using a wrapping key. To maintain security, this encryption key must occasionally be changed. For this reason, a command is provided on the firewall to rotate the wrapping key which changes the master key encryption. The frequency of this wrapping key rotation depends on your application.

Master key encryption using an HSM is not supported on firewalls configured in FIPS of CC mode.

The following topics describe how to encrypt the master key initially and how to refresh the master key encryption:

- Encrypt the Master Key
- Refresh the Master Key Encryption

## Encrypt the Master Key

If you have not previously encrypted the master key on a device, use the following procedure to encrypt it. Use this procedure for first time encryption of a key, or if you define a new master key and you want to encrypt it. If you want to refresh the encryption on a previously encrypted key, see Refresh the Master Key Encryption.

| Encrypt a Master Key Using an HSM | |
|---|---|
| 1. | Select **Device > Master Key and Diagnostics.** |
| Step 8 | Specify the key that is currently used to encrypt all of the private keys and passwords on the firewall in the **Master Key** field. |
| Step 9 | If changing the master key, enter the new master key and confirm. |
| Step 10 | Select the **HSM** check box.<br>**Life Time:** The number of days and hours after which the master key expires (range 1-730 days).<br>**Time for Reminder:** The number of days and hours before expiration when the user is notified of the impending expiration (range 1-365 days). |
| Step 11 | Click **OK**. |

## Refresh the Master Key Encryption

As a best practice, refresh the master key encryption on a regular basis by rotating the master key wrapping key on the HSM. This command is the same for both the SafeNet Luna SA and Thales Nshield Connect HSMs.

 © Palo Alto Networks, Inc.

**Refresh the Master Key Encryption**

1. Use the following CLI command to rotate the wrapping key for the master key on an HSM:

   ```
   > request hsm mkey-wrapping-key-rotation
   ```

   If the master key is encrypted on the HSM, the CLI command will generate a new wrapping key on the HSM and encrypt the master key with the new wrapping key.

   If the master key is not encrypted on the HSM, the CLI command will generate new wrapping key on the HSM for future use.

   The old wrapping key is not deleted by this command.

## Store Private Keys on an HSM

For added security, the private keys used to enable SSL/TLS decryption—both SSL forward proxy and SSL inbound inspection—can be secured with an HSM as follows:

- **SSL forward proxy**—The private key in the CA certificate that is used to sign certificates in SSL/TLS forward proxy operations can be stored on the HSM. The firewall will then send the certificates it generates during SSL/TLS forward proxy operations to the HSM for signing before forwarding them on to the client.
- **SSL inbound inspection**—The private keys for the internal servers for which you are doing SSL/TLS inbound inspection can be stored on the HSM.

For instructions on importing the private keys onto the HSM, refer to the documentation from your HSM provider. After the required keys are on the HSM, you can configure the firewall to locate the keys as follows:

| **Store Private Keys on an HSM** | | |
|---|---|---|
| Step 1 | Import the private keys used in your SSL forward proxy and/or SSL inbound inspection deployments onto the HSM. | For instructions on importing the private keys onto the HSM, refer to the documentation from your HSM provider. |
| Step 2 | (Thales Nshield Connect only) Sync the key data from the HSM remote file system to the firewall. | 1. From the firewall web interface, select **Device > Setup > HSM**.<br>2. Select **Synchronize with Remote Filesystem** in the Hardware Security Operations section. |
| Step 3 | Import the certificate(s) that correspond to the private key(s) you are storing on the HSM onto the firewall. | 1. From the firewall web interface, select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Click **Import**.<br>3. Enter the **Certificate Name**.<br>4. Enter filename of the **Certificate File** you imported to the HSM.<br>5. Select the appropriate **File Format** from the drop-down.<br>6. Select the **Private Key resides on Hardware Security Module** check box.<br>7. Click **OK** and **Commit**. |
| Step 4 | (Forward trust certificates only) Enable the certificate for use in SSL/TLS Forward Proxy. | 1. Select the **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Locate the certificate you imported in Step 3.<br>3. Select the **Forward Trust Certificate** check box.<br>4. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

| Store Private Keys on an HSM (Continued) | | |
|---|---|---|
| Step 5 | Verify that the certificate has been successfully imported to the firewall. | 1. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Locate the certificate you imported in Step 3.<br>3. In the Key column notice the following:<br>If a Lock icon is displayed, the private key for the certificate can be found on the HSM.<br>If an Error icon is displayed, the private key is not imported to the HSM or the HSM is not properly authenticated or connected. |

## Manage the HSM Deployment

| Manage HSM | |
|---|---|
| • View the HSM configuration settings. | Select **Device > Setup > HSM**. |
| • Display detailed HSM information. | Select **Show Detailed Information** from the Hardware Security Operations section.<br>Information regarding the HSM servers, HSM HA status, and HSM hardware is displayed. |
| • Export Support file | Select **Export Support File** from the Hardware Security Operations section.<br>A test file is created to help customer support when addressing a problem with an HSM configuration on the firewall. |
| • Reset HSM configuration. | Select **Reset HSM Configuration** from the Hardware Security Operations section.<br>Selecting this option removes all HSM connections. All authentication procedures must be repeated after using this option. |

 © Palo Alto Networks, Inc.

# High Availability

High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in a two-device cluster provides redundancy and allows you to ensure business continuity.

The Palo Alto Networks firewalls support stateful active/passive or active/active high availability with session and configuration synchronization. Some models of the firewall, such as the VM-Series firewall and the PA-200, only support HA lite without session synchronization capability. The following topics provide more information about high availability and how to configure it in your environment.

- HA Overview
- HA Concepts
- Set Up Active/Passive HA
- HA Resources

# HA Overview

On Palo Alto Networks firewalls, you can set up two devices as an HA pair. HA allows you to minimize downtime by making sure that an alternate device is available in the event that the primary device fails. The devices use dedicated or in-band HA ports on the firewall to synchronize data—network, object, and policy configurations—and to maintain state information. Device specific configuration such as management port IP address or administrator profiles, HA specific configuration, log data, and the Application Command Center (ACC) information is not shared between devices. For a consolidated application and log view across the HA pair, you must use Panorama, the Palo Alto Networks centralized management system.

When a failure occurs on the active device and the passive device takes over the task of securing traffic, the event is called a failover. The conditions that trigger a failover are:

- One or more of the monitored interfaces fail. (Link Monitoring)
- One or more of the destinations specified on the device cannot be reached. (Path Monitoring)
- The device does not respond to heartbeat polls. (Heartbeat Polling and Hello messages)

After you understand the HA Concepts, continue to Set Up Active/Passive HA.

 © Palo Alto Networks, Inc.

# HA Concepts

The following topics provide conceptual information about how HA works on a Palo Alto Networks firewall:


- HA Modes
- HA Links and Backup Links
- Device Priority and Preemption
- Failover Triggers
- HA Timers


## HA Modes

You can set up the firewalls for HA in two modes:

- **Active/Passive**— One device actively manages traffic while the other is synchronized and ready to transition to the active state, should a failure occur. In this configuration, both devices share the same configuration settings, and one actively manages traffic until a path, link, system, or network failure occurs. When the active device fails, the passive device takes over seamlessly and enforces the same policies to maintain network security. Active/passive HA is supported in the virtual wire, Layer 2 and Layer 3 deployments. For information on setting up your devices in an active/passive configuration, see Configure Active/Passive HA.

The PA-200 and the VM-Series firewalls support a lite version of active/passive HA. HA lite provides configuration synchronization and some runtime data synchronization such as IPSec security associations. It does not support any session synchronization, and therefore, HA Lite does not offer stateful failover.

- **Active/Active**— Both the devices in the pair are active and processing traffic, and work synchronously to handle session setup and session ownership. The active/active deployment is supported in virtual wire and Layer 3 deployments, and is only recommended for networks with asymmetric routing. For information on setting up the devices in an active/active configuration, refer to the Active/Active High Availability Tech Note.

## HA Links and Backup Links

The devices in an HA pair use HA links to synchronize data and maintain state information. Some models of the firewall have dedicated HA ports—Control link (HA1) and Data link (HA2), while others require you to use the in-band ports as HA links.

On devices with dedicated HA ports such as the PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7050 firewalls (see HA Ports on the PA-7050 Firewall), use the dedicated HA ports to manage communication and synchronization between the devices. For devices without dedicated HA ports such as the PA-200, PA-500, and PA-2000 Series firewalls, as a best practice use the management port for the HA1 link to allow for a direct connection between the management planes on the devices, and an in-band port for the HA2 link.

The HA1 and HA2 links provide synchronization for functions that reside on the management plane. Using the dedicated HA interfaces on the management plane is more efficient than using the in-band ports as this eliminates the need to pass the synchronization packets over the dataplane.

- **Control Link**: The HA1 link is used to exchange hellos, heartbeats, and HA state information, and management plane sync for routing, and User-ID information. This link is also used to synchronize configuration changes on either the active or passive device with its peer. The HA1 link is a Layer 3 link and requires an IP address.
  **Ports used for HA1**: TCP port 28769 and 28260 for clear text communication; port 28 for encrypted communication (SSH over TCP).

- **Data Link**: The HA2 link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between devices in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active device to the passive device. The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.
  **Ports used for HA2:** The HA data link can be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets.
  Additionally, an HA3 link is used in Active/Active HA deployments. When there is an asymmetric route, the HA3 link is used for forwarding packets to the HA peer that owns the session. The HA3 link is a Layer 2 link and it does not support Layer 3 addressing or encryption.

- **Backup Links**: Provide redundancy for the HA1 and the HA2 links. In-band ports are used as backup links for both HA1 and HA2. Consider the following guidelines when configuring backup HA links:
  - The IP addresses of the primary and backup HA links must not overlap each other.
  - HA backup links must be on a different subnet than the primary HA links.
  - HA1-backup and HA2-backup ports must be configured on separate physical ports. The HA1-backup link uses port 28770 and 28260.

Palo Alto Networks recommends enabling heartbeat backup (uses port 28771 on the MGT interface) if you use an in-band port for the HA1 or the HA1 backup links.

## HA Ports on the PA-7050 Firewall

For HA connectivity on the PA-7050, refer to the following table for details on which ports on the Switch Management Card (SMC) are mandated and where ports on the Network Processing Card (NPC) are suitable. For an overview of the Modules and Interface cards on the PA-7050 firewall, refer to the PA-7050 Hardware Reference Guide.

The following ports on the SMC are designed for HA connectivity:

 © Palo Alto Networks, Inc.

| HA Links and Backup Links | Ports on the SMC | Description |
|---|---|---|
| **Control Link** | HA1-A<br>Speed: Ethernet 10/100/1000 | Used for HA control and synchronization. Connect this port directly from the HA1-A port on the first device to the HA1-A on the second device in the pair, or connect them together through a switch or router.<br>HA1 cannot be configured on NPC data ports or the MGT port. |
| **Control Link Backup** | HA1-B<br>Speed: Ethernet 10/100/1000 port | Used for HA control and synchronization as a backup for HA1-A. Connect this port directly from the HA1-B port on the first device to the HA1-B on the second device in the pair, or connect them together through a switch or router.<br>HA1 Backup cannot be configured on NPC data ports or the MGT port. |
| **Data Link** | HSCI-A<br>(High Speed Chassis Interconnect) | Quad Port SFP (QSFP) interfaces used to connect two PA-7050 firewalls in an HA configuration. Each port is comprised of four 10 gigabit links internally for a combined speed of 40 gigabits and is used for HA2 data link in an active/passive configuration. When in active/active mode, the port is also used for HA3 packet forwarding for asymmetrically routed sessions that require Layer 7 inspection for App-ID and Content-ID.<br>In a typical installation, HSCI-A on the first chassis connects directly to HSCI-A on the second chassis and HSCI-B on the first chassis connects to HSCI-B on the second chassis. This will provide full 80 gigabit transfer rates. In software, both ports (HSCI-A and HSCI-B) are treated as one HA interface.<br>The HSCI ports are not routable and must be connected directly to each other.<br>Palo Alto Networks recommends using the dedicated HSCI ports for both HA2 and HA3 connections. However, the HA2 and HA3 links can be configured on NPC data ports, if needed. |
| **Data Link Backup** | HSCI-B<br>(High Speed Chassis Interconnect) | The Quad Port SFP (QSFP) interfaces (see description above) in the HSCI-B port is used to increase the bandwidth for HA2/HA3 purposes.<br>The HSCI ports are not routable and must be connected directly to each other.<br>Palo Alto Networks recommends using the dedicated HSCI-B ports for both HA2 and HA3 backup connections. The HA2/HA3 backup link can be configured on the NPC data ports, if needed. |

## Device Priority and Preemption

The devices in an HA pair can be assigned a *device priority* value to indicate a preference for which device should assume the active role and manage traffic. If you need to use a specific device in the HA pair for actively securing traffic, you must enable the preemptive behavior on both the firewalls and assign a device priority value for each

device. The device with the lower numerical value, and therefore *higher priority*, is designated as active and manages all traffic on the network. The other device is in a passive state, and synchronizes configuration and state information with the active device so that it is ready to transition to an active state should a failure occur.

By default, preemption is disabled on the firewalls and must be enabled on both devices. When enabled, the preemptive behavior allows the firewall with the *higher priority* (lower numerical value) to resume as active after it recovers from a failure. When preemption occurs, the event is logged in the system logs.

## Failover Triggers

When a failure occurs on the active device and the passive device takes over the task of securing traffic, the event is called a failover. A failover is triggered when a monitored metric on the active device fails. The metrics that are monitored for detecting a device failure are:

- **Heartbeat Polling and Hello messages**

  The firewalls use hello message and heartbeats to verify that the peer device is responsive and operational. Hello messages are sent from one peer to the other at the configured *Hello Interval* to verify the state of the device. The heartbeat is an ICMP ping to the HA peer over the control link, and the peer responds to the ping to establish that the devices are connected and responsive. By default, the interval for the heartbeat is 1000 milliseconds. For details on the HA timers that trigger a failover, see HA Timers.

- **Link Monitoring**

  The physical interfaces to be monitored are grouped into a link group and their state (link up or link down) is monitored. A link group can contain one or more physical interfaces. A device failure is triggered when any or all of the interfaces in the group fail. The default behavior is failure of any one link in the link group will cause the device to change the HA state to non-functional to indicate a failure of a monitored object.

- **Path Monitoring**

  Monitors the full path through the network to mission-critical IP addresses. ICMP pings are used to verify reachability of the IP address. The default interval for pings is 200ms. An IP address is considered unreachable when 10 consecutive pings (the default value) fail, and a device failure is triggered when any or all of the IP addresses monitored become unreachable. The default behavior is any one of the IP addresses becoming unreachable will cause the device to change the HA state to non-functional to indicate a failure of a monitored object.

In addition to the failover triggers listed above, a failover also occurs when the administrator places the device is a suspended state or if preemption occurs.

On the PA-3000 Series, PA-5000 Series, and PA-7050 firewalls, a failover can occur when an internal health check fails. This health check is not configurable and is enabled to verify the operational status for all the components within the firewall.

 © Palo Alto Networks, Inc.

## HA Timers

High Availability (HA) timers are used to detect a firewall failure and trigger a failover. To reduce the complexity in configuring HA timers, you can select from three profiles have been added: **Recommended**, **Aggressive** and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.

Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

The following table describes each timer included in the profiles and the current preset values across the different hardware models; these values are for current reference only and can change in a subsequent release.

**Recommended/Aggressive HA Timer Values by Platform**

| Timers | Description | PA-7050<br>PA-5000 Series<br>PA-4000 Series<br>PA-3000 Series | PA-2000 Series<br>PA-500 Series<br>PA-200 Series<br>VM-Series | Panorama VM<br>M-100 |
|---|---|---|---|---|
| Monitor fail hold up time | The interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is recommended to avoid an HA failover due to the occasional flapping of neighboring devices. | 0/0 | 0/0 | 0/0 |
| Preemption hold time | Time a passive or active-secondary device will wait before taking over as the active or active-primary device. | 1/1 | 1/1 | 1/1 |
| Heartbeat interval | The frequency at which the HA peers exchange heartbeat messages in the form of an ICMP ping. | 1000/1000 | 2000/1000 | 2000/1000 |
| Promotion hold time | Time that the passive device (in active/passive mode) or the active-secondary device (in active/active mode) will wait before taking over as the active or active-primary device after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made. | 2000/500 | 2000/500 | 2000/500 |

| Timers | Description | PA-7050<br>PA-5000 Series<br>PA-4000 Series<br>PA-3000 Series | PA-2000 Series<br>PA-500 Series<br>PA-200 Series<br>VM-Series | Panorama VM<br>M-100 |
|---|---|---|---|---|
| Additional master hold up time | This time interval is applied to the same event as Monitor Fail Hold Up Time (range 0-60000 ms, default 500 ms). The additional time interval is applied only to the active device in active/passive mode and to the active-primary device in active/active mode. This timer is recommended to avoid a failover when both devices experience the same link/path monitor failure simultaneously. | 500/500 | 500/500 | 7000/5000 |
| Hello interval | The time interval in milliseconds between the hello packets that are sent to verify that the HA functionality on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms. | 8000/8000 | 8000/8000 | 8000/8000 |
| Maximum no. of flaps | A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. This value indicates the maximum number of flaps that are permitted before the firewall is determined to be suspended and the passive firewall takes over (range 0-16, default 3). | 3/3 | 3/3 | Not Applicable |

 © Palo Alto Networks, Inc.

# Set Up Active/Passive HA


- Prerequisites for Active/Passive HA
- Configuration Guidelines for Active/Passive HA
- Configure Active/Passive HA
- Define HA Failover Conditions
- Verify Failover


© Palo Alto Networks, Inc.

## Prerequisites for Active/Passive HA

To set up high availability on your Palo Alto Networks firewalls, you need a pair of firewalls that meet the following requirements:

- **The same model**—both the devices in the pair must be of the same hardware model or virtual machine model.
- **The same PAN-OS version**—both the devices should be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases. They must also both have the same multiple virtual systems capability (single or multi vsys).
- **The same type of interfaces**—dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
  - Determine the IP address for the HA1 (control) connection between the device pair. The HA1 IP address for both peers must be on the same subnet if they are directly connected or are connected to the same switch.

    For devices without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both devices. However, because the management ports will not be directly cabled between the devices, make sure that you have a route that connects these two interfaces across your network.
  - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use Layer 3 *only* if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
- **The same set of licenses**—Licenses are unique to each device and cannot be shared between the devices. Therefore, you must license both devices identically. If both devices do not have an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.

If you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration, it is recommended that you perform a factory reset on the new firewall. This will ensure that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary device to the newly introduced device with the clean config.

 © Palo Alto Networks, Inc.

## Configuration Guidelines for Active/Passive HA

To set up an active (PeerA) passive (PeerB) pair in HA, you must configure some options identically on both devices and some independently (non-matching) on each device. These HA settings are not synchronized between the devices. For details on what is/is not synchronized, refer to HA Synchronization.

To proceed with the instructions on configuring the devices in HA, see Configure Active/Passive HA.

The following table lists the settings that you must configure identically on both devices:

| Identical Configuration Settings on PeerA and PeerB |
|---|
| • HA must be enabled on both devices.<br>• Both device must have the same Group ID value. The Group ID value is used to create a virtual MAC address for all the configured interfaces. The format of the virtual MAC is 00-1B-17:00: xx: yy where<br>00-1B-17: vendor ID; 00: fixed; xx: HA group ID; yy: interface ID.<br>When a new active device takes over, Gratuitous ARPs are sent from each of the connected interfaces of the new active member to inform the connected Layer 2 switches of the virtual MAC address' new location.<br>• If using in-band ports, the interfaces for the HA1 and HA2 links must be set to type HA.<br>• The HA mode must be set to Active Passive.<br>• If required, preemption must be enabled on both devices. The device priority value, however, must not be identical.<br>• If required, encryption on the HA1 link (for communication between the HA peers) must be configured on both devices.<br>• Based on the combination of HA1 and HA1 Backup ports you are using, use the following recommendations to decide whether you should enable heartbeat backup:<br>  • HA1: Dedicated HA1 port<br>HA1 Backup: In-band port<br>Recommendation: Enable Heartbeat Backup<br>  • HA1: Dedicated HA1 port<br>HA1 Backup: Management port<br>Recommendation: Do not enable Heartbeat Backup<br>  • HA1: In-band port<br>HA1 Backup: In-band port<br>Recommendation: Enable Heartbeat Backup<br>  • HA1: Management port<br>HA1 Backup: In-band port<br>Recommendation: Do not enable Heartbeat Backup |

The following table lists the settings that must be configured independently on each device:

| Independent Configuration Settings | PeerA | PeerB |
|---|---|---|
| Control Link | IP address of the HA1 link configured on this device (PeerA). | IP address of the HA1 link configured on this device (PeerB). |
| | For devices without dedicated HA ports, use the management port IP address for the control link. | |
| Data Link<br>The data link information is synchronized between the devices after HA is enabled and the control link is established between the devices. | By default, the HA2 link uses Ethernet/Layer 2.<br>If using a Layer 3 connection, configure the IP address for the data link on this device (PeerA). | By default, the HA2 link uses Ethernet/Layer 2.<br>If using a Layer 3 connection, configure the IP address for the data link on this device (PeerB). |
| Device Priority (required, if preemption is enabled) | The device you plan to make active must have a lower numerical value than its peer. So, if Peer A is to function as the active device, keep the default value of 100 and increment the value on PeerB. | If PeerB is passive, set the device priority value to a number larger than that on PeerA. For example, set the value to 110. |
| Link Monitoring—Monitor one or more physical interfaces that handle vital traffic on this device and define the failure condition. | Select the physical interfaces on the firewall that you would like to monitor and define the failure condition (all or any) to trigger a failover. | Pick a similar set of physical interfaces that you would like to monitor on this firewall and define the failure condition (all or any) to trigger a failover. |
| Path Monitoring—Monitor one or more destination IP addresses that the firewall can use ICMP pings to ascertain responsiveness. | Define the failure condition (all or any), ping interval and the ping count. This is particularly useful for monitoring the availability of other interconnected networking devices. For example, monitor the availability of a router that connects to a server, connectivity to the server itself, or some other vital device that is in the flow of traffic.<br>Make sure that the node/device that you are monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a a path monitoring failure and trigger a failover. | Pick a similar set of devices or destination IP addresses that can be monitored for determining the failover trigger for PeerB. Define the failure condition (all or any), ping interval and the ping count. |

 © Palo Alto Networks, Inc.

## Configure Active/Passive HA

The following procedure shows how to configure a pair of firewalls in an active/passive deployment as depicted in the following example topology.

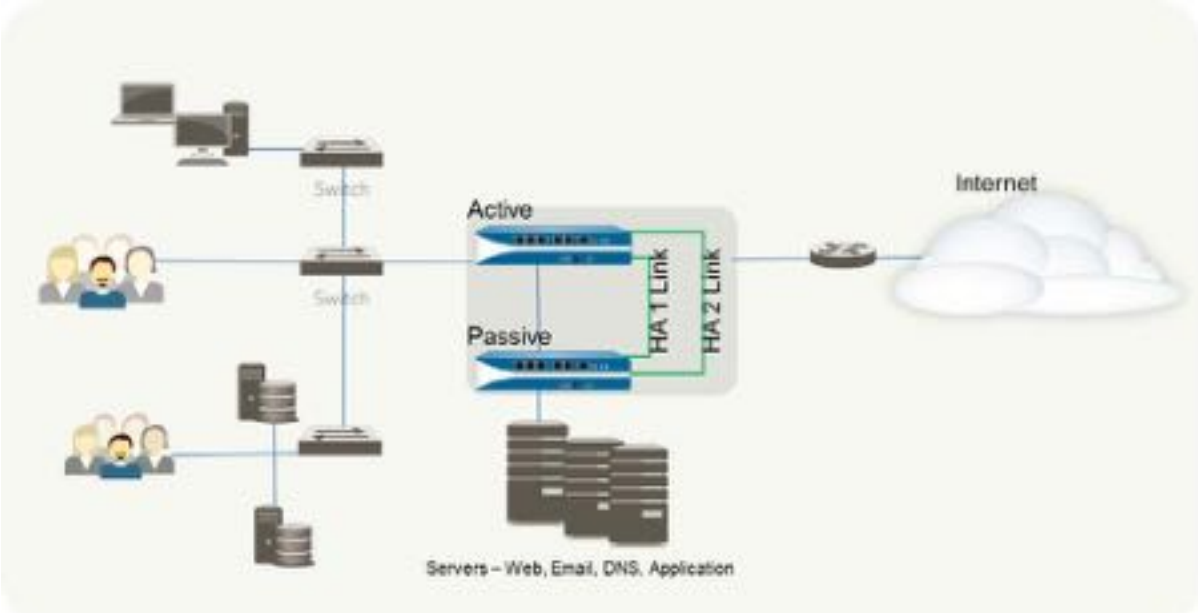


| Connect and Configure the Devices | |
|---|---|
| Step 1 Connect the HA ports to set up a physical connection between the devices. | • For devices with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on the device pair. Use a crossover cable if the devices are directly connected to each other.<br>• For devices without dedicated HA ports, select two data interfaces for the HA2 link and the backup HA1 link. Then, use an Ethernet cable to connect these in-band HA interfaces across both devices. Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network. |
| Pick a device in the pair and complete these tasks: | |
| Step 2 Enable ping on the management port. Enabling ping allows the management port to exchange heartbeat backup information. | 1. Select **Device > Setup > Management** and then click the Edit icon in the Management Interface Settings section of the screen.<br>2. Select **Ping** as a service that is permitted on the interface. |
| Step 3 If the device does not have dedicated HA ports, set up the data ports to function as HA ports.<br>For devices with dedicated HA ports continue to Step 4. | 1. Select **Network > Interfaces.**<br>2. Confirm that the link is up on the ports that you want to use.<br>3. Select the interface and set the interface type to HA.<br> 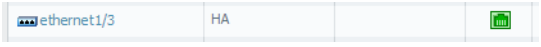 <br>4. Set the **Link Speed** and **Link Duplex** settings, as appropriate. |

| Connect and Configure the Devices (Continued) | | |
|---|---|---|
| Step 4 | Set up the control link connection.<br><br>This example shows an in-band port that is set to interface type HA.<br><br>For devices that use the management port as the control link, the IP address information is automatically pre-populated. | 1. In **Device > High Availability > General**, edit the **Control Link (HA1)** section.<br>2. Select the interface that you have cabled for use as the HA1 link in the **Port** drop down menu. Set the IP address and netmask. Enter a Gateway IP address only if the HA1 interfaces are on separate subnets. Do not add a gateway if the devices are directly connected.<br> |
| Step 5 | (Optional) Enable encryption for the control link connection.<br><br>This is typically used to secure the link if the two devices are not directly connected, that is if the ports are connected to a switch or a router. | 1. Export the HA key from a device and import it into the peer device.<br>a. Select **Device > Certificate Management > Certificates.**<br>b. Select **Export HA key.** Save the HA key to a network location that the peer device can access.<br>c. On the peer device, navigate to **Device > Certificate Management > Certificates,** and select **Import HA key** to browse to the location that you saved the key and import it in to the peer device.<br>2. Select **Device > High Availability > General**, edit the **Control Link (HA1)** section.<br>3. Select **Encryption Enabled.** |
| Step 6 | Set up the backup control link connection. | 1. In **Device > High Availability > General**, edit the **Control Link (HA1 Backup)** section.<br>2. Select the HA1 backup interface and set the IP address and netmask.<br> |

 © Palo Alto Networks, Inc.

| Connect and Configure the Devices (Continued) | | |
|---|---|---|
| Step 7 | Set up the data link connection (HA2) and the backup HA2 connection between the devices. | 1. In **Device > High Availability > General**, edit the **Data Link (HA2)** section.<br>2. Select the interface for the data link connection.<br>3. Select the **Transport** method. The default is **ethernet**, and will work when the HA pair is connected directly or through a switch. If you need to route the data link traffic through the network, select **IP** or **UDP** as the transport mode.<br>4. If you use IP or UDP as the transport method, enter the IP address and netmask.<br> 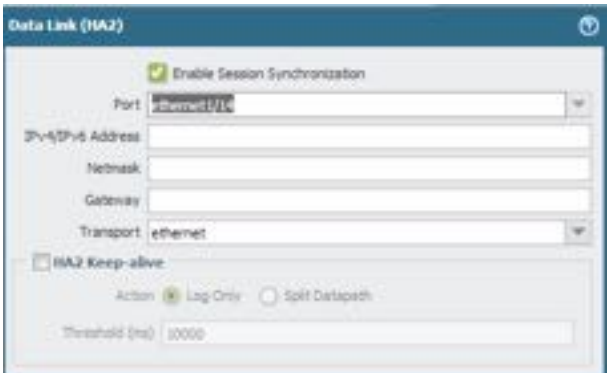 <br>5. Verify that **Enable Session Synchronization** is selected.<br>6. Select **HA2 Keep-alive** to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set (default is 10000 ms), the defined action will occur. For active/passive configuration, a critical system log message is generated when an HA2 keep-alive failure occurs.<br>You can configure the HA2 keep-alive option on both devices, or just one device in the HA pair. If the option is only enabled on one device, only that device will send the keep-alive messages. The other device will be notified if a failure occurs.<br>7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the IP address and netmask. |
| Step 8 | Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port.<br><br>You do not need to enable heartbeat backup if you are using the management port for the control link. | 1. In **Device > High Availability > General**, edit the **Election Settings** section.<br>2. Select **Heartbeat Backup**.<br>The heartbeat backup link is used for transmitting redundant heartbeats and hello messages. To allow the heartbeats to be transmitted between the devices, you must verify that the management port across both peers can route to each other. |

| Connect and Configure the Devices (Continued) | |
|---|---|
| Step 9 Set the device priority and enable preemption.<br><br>This setting is only required if you wish to make sure that a specific device is the preferred active device. For information, see Device Priority and Preemption. | 1. In **Device > High Availability > General**, edit the **Election Settings** section.<br>2. Set the numerical value in **Device Priority**. Make sure to set a lower numerical value on the device that you want to assign a higher priority to.<br><br>If both firewalls have the same device priority value, the firewall with the lowest MAC address on the HA1 control link will become the active device.<br>3. Select **Preemptive**.<br><br>You must enable preemptive on both the active and the passive device. |
| Step 10 (Optional) Modify the failover timers.<br><br>By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments. | 1. In **Device > High Availability > General**, edit the **Election Settings** section.<br>2. Select the **Aggressive** profile for triggering failover faster; select **Advanced** to define custom values for triggering failover in your set up.<br><br>To view the preset value for an individual timer included in a profile, select **Advanced** and click **Load Recommended** or **Load Aggressive**. The preset values for your hardware model will be displayed on screen. |
| Step 11 (Optional, only configured on the passive device) Modify the link status of the HA ports on the passive device.<br><br>The passive link state is **shutdown**, by default. After you enable HA, the link state for the HA ports on the active device will be green and those on the passive device will be down and display as red. | Setting the link state to **Auto** allows for reducing the amount of time it takes for the passive device to take over when a failover occurs and it allows you to monitor the link state.<br><br>To enable the link status on the passive device to stay up and reflect the cabling status on the physical interface:<br>1. In **Device > High Availability > General**, edit the **Active Passive Settings** section.<br>2. Set the **Passive Link State** to **Auto**.<br><br>The auto option decreases the amount of time it takes for the passive device to take over when a failover occurs.<br><br>Although the interface displays green (as cabled and up) it continues to discard all traffic until a failover is triggered.<br><br>When you modify the passive link state, make sure that the adjacent devices do not forward traffic to the passive firewall based only on the link status of the device. |

 © Palo Alto Networks, Inc.

| Connect and Configure the Devices (Continued) | |
|---|---|
| Step 12 Enable HA. | 1. Select **Device > High Availability > General,** edit the **Setup** section.<br>2. Select **Enable HA.**<br>3. Set a **Group ID**. This ID uniquely identifies each HA pair on your network, and is essential if you have multiple HA pairs that share the same broadcast domain on your network.<br>4. Set the mode to **Active Passive**.<br>5. Select **Enable Config Sync.** This setting enables the synchronization of the configuration settings between the active and the passive device.<br>6. Enter the IP address assigned to the control link of the peer device in **Peer HA1 IP Address**.<br> 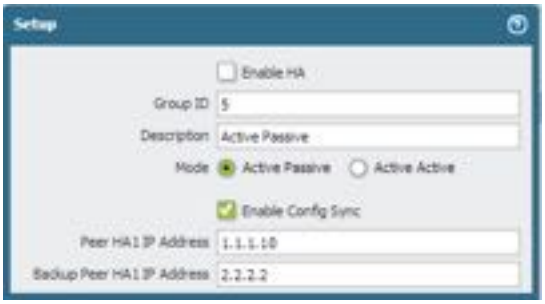 <br>For devices without dedicated HA ports, if the peer uses the management port for the HA1 link, enter the management port IP address of the peer.<br>7. Enter the **Backup HA1 IP Address**. |
| Step 13 Save your configuration changes. | Click **Commit**. |
| Step 14 Complete Step 2 through Step 13 on the other device in the HA pair. | |
| Step 15 After you finish configuring both devices, verify that the devices are paired in active/passive HA. | 1. Access the **Dashboard** on both devices, and view the **High Availability** widget.<br>2. On the active device, click the **Sync to peer** link.<br>3. Confirm that the devices are paired and synced, as shown below: |

**Connect and Configure the Devices (Continued)**

On the passive device: The state of the local device should display **passive** and the configuration is **synchronized.**

On the active device: The state of the local device should display **active** and the configuration is **synchronized**.

 © Palo Alto Networks, Inc.

# Define HA Failover Conditions

| Configure the Failover Triggers | | |
|---|---|---|
| Step 1 | To configure link monitoring, define the interfaces that you would like to monitor. A change in the link state of these interface will trigger a failover. | 1. Select **Device > High Availability > Link and Path Monitoring.**<br>2. In the **Link Group** section, click **Add.**<br>3. Name the **Link Group, Add** the interfaces to monitor, and select the **Failure Condition** for the group. The Link group you define is added to the **Link Group** section.<br> |
| Step 2 | (Optional) Modify the failure condition for the Link Groups that you configured (in the preceding step) on the device.<br><br>By default, the device will trigger a failover when any monitored link fails. | 1. Select the **Link Monitoring** section.<br>2. Set the **Failure Condition** to **All.**<br>The default setting is **Any**. |
| Step 3 | To configure path monitoring, define the destination IP addresses that the firewall should ping to verify network connectivity. | 1. In the **Path Group** section of the **Device > High Availability > Link and Path Monitoring** tab, pick the **Add option for your set up:** Virtual Wire, VLAN, or Virtual Router.<br>2. Select the appropriate item from the drop-down list for the **Name** and **Add** the IP addresses (source and/or destination, as prompted) that you wish to monitor. Then select the **Failure Condition** for the group. The path group you define is added to the **Path Group** section.<br> |
| Step 4 | (Optional) Modify the failure condition for all Path Groups configured on the device.<br><br>By default, the device will trigger a failover when any monitored path fails. | Set the **Failure Condition** to **All.**<br><br>The default setting is **Any**. |
| Step 5 | Save your changes. | Click **Commit**. |

If you are using SNMPv3 to monitor the firewalls, note that the SNMPv3 Engine ID is unique to each device; the EngineID is not synchronized between the HA pair and therefore, allows you to independently monitor each device in the HA pair. For information on setting up SNMP, see Set Up SNMP Trap Destinations.

Because the EngineID is generated using the device's unique serial number, on the VM-Series firewall you must apply a valid license in order to obtain a unique EngineID for each firewall.

## Verify Failover

To test that your HA configuration works properly trigger a manual failover and verify that the devices transition states successfully.

| Verify Failover | |
|---|---|
| Step 1 Suspend the active device. | Click the **Suspend local device** link on the **Device > High Availability > Operational Commands** tab.<br> 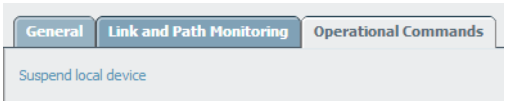 |
| Step 2 Verify that the passive device has taken over as active. | On the **Dashboard**, verify that the state of the passive device changes to **active** in the **High Availability** widget.<br> |
| Step 3 Restore the suspended device to a functional state. Wait for a couple minutes, and then verify that preemption has occurred, if preemptive is enabled. | 1. On the device you previously suspended, select the **Make local device functional** link on the **Device > High Availability > Operational Commands** tab.<br> 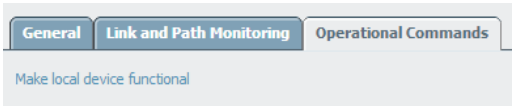 <br>2. In the **High Availability** widget on the **Dashboard,** confirm that the device has taken over as the active device and that the peer is now in a passive state.<br> |

 © Palo Alto Networks, Inc.

# HA Resources

For more information on HA, refer to the following sources:

- Active/Active HA
- High Availability Synchronization
- High Availability Failover Optimization
- Upgrading an HA pair
- Examples: Deploying HA

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

# Reports and Logging

The firewall provides reports and logs that are useful for monitoring activity on your network. You can monitor the logs and filter the information to generate reports with predefined or customized views. You can for example, use the predefined templates to generate reports on a user's activity or analyze the reports and logs to interpret unusual behavior on your network and generate a custom report on the traffic pattern. The following topics describe how to view, manage, customize, and generate the reports and logs on the firewall:

- Use the Dashboard
- Use the Application Command Center
- Use App-Scope
- Take Packet Captures
- Monitor the Firewall
- Forward Logs to External Services
- Monitor the Firewall Using SNMP
- Manage Reporting
- Identify Firewall Interfaces in External Monitoring Systems
- Manage Reporting
- Syslog Field Descriptions

# Use the Dashboard

The **Dashboard** tab widgets show general device information, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available widgets are displayed by default, but each administrator can remove and add individual widgets, as needed.

Click the refresh icon to update the Dashboard or an individual widget. To change the automatic refresh interval, select an interval from the drop-down (**1 min**, **2 mins**, **5 mins**, or **Manual**). To add a widget to the Dashboard, click the widget drop-down, select a category and then the widget name. To delete a widget, click in the title bar.

The following table describes the Dashboard widgets.

| Dashboard Charts | Descriptions |
|---|---|
| Top Applications | Displays the applications with the most sessions. The block size indicates the relative number of sessions (mouse-over the block to view the number), and the color indicates the security risk—from green (lowest) to red (highest). Click an application to view its application profile. |
| Top High Risk Applications | Similar to Top Applications, except that it displays the highest-risk applications with the most sessions. |
| General Information | Displays the device name, model, PAN-OS software version, the application, threat, and URL filtering definition versions, the current date and time, and the length of time since the last restart. |
| Interface Status | Indicates whether each interface is up (green), down (red), or in an unknown state (gray). |
| Threat Logs | Displays the threat ID, application, and date and time for the last 10 entries in the Threat log. The threat ID is a malware description or URL that violates the URL filtering profile. |
| Config Logs | Displays the administrator username, client (Web or CLI), and date and time for the last 10 entries in the Configuration log. |
| Data Filtering Logs | Displays the description and date and time for the last 60 minutes in the Data Filtering log. |
| URL Filtering Logs | Displays the description and date and time for the last 60 minutes in the URL Filtering log. |
| System Logs | Displays the description and date and time for the last 10 entries in the System log.<br>A **`Config installed`** entry indicates configuration changes were committed successfully. |
| System Resources | Displays the Management CPU usage, Data Plane usage, and the Session Count, which displays the number of sessions established through the firewall. |
| Logged In Admins | Displays the source IP address, session type (Web or CLI), and session start time for each administrator who is currently logged in. |
| ACC Risk Factor | Displays the average risk factor (1 to 5) for the network traffic processed over the past week. Higher values indicate higher risk. |

 © Palo Alto Networks, Inc.

| Dashboard Charts | Descriptions |
|---|---|
| High Availability | If high availability (HA) is enabled, indicates the HA status of the local and peer device—green (active), yellow (passive), or black (other). For more information about HA, see High Availability. |
| Locks | Shows configuration locks taken by administrators. |

# Use the Application Command Center

The **ACC** tab visually depicts trends and historic view of traffic on your network.


- ACC Risk Level
- ACC Charts
- ACC Detail Pages
- Use the ACC


 © Palo Alto Networks, Inc.

## ACC Risk Level

The **ACC** tab displays the overall risk level for all network traffic, the risk levels and number of threats detected for the most active and highest-risk applications on your network, and the number of threats detected from the busiest application categories and from all applications at each risk level. Use the ACC to view application data for the past hour, day, week, month, or any custom-defined time frame. **Risk** levels (1=lowest to 5=highest) indicate the application's relative security risk based on criteria such as whether the application can share files, is prone to misuse, or tries to evade firewalls.

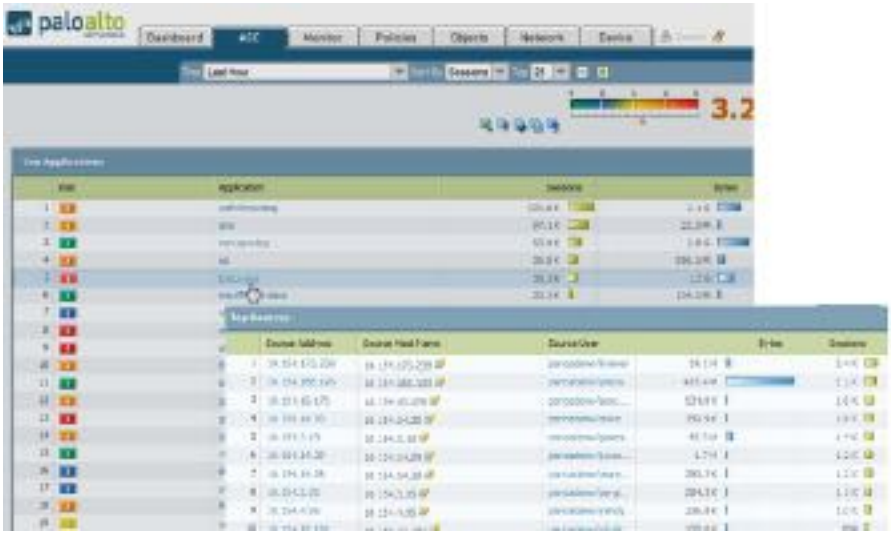

## ACC Charts

There are five charts displayed on the **Application Command Center (ACC)** tab:

- Application
- URL Filtering
- Threat Prevention
- Data Filtering
- HIP Matches

The following table describes the charts displayed on the **ACC** tab:

| ACC Chart | Description |
|---|---|
| Application | Displays application information grouped by the following attributes:<br>• Applications<br>• High risk applications<br>• Categories<br>• Sub Categories<br>• Technology<br>• Risk<br>Each chart can include the number of sessions, bytes transmitted and received, number of threats, application category, application subcategories, application technology, and risk level, as applicable. |
| URL Filtering | Displays URL/category information grouped by the following attributes:<br>• URL Categories<br>• URLs<br>• Blocked URL Categories<br>• Blocked URLs<br>Each chart can include the URL, URL category, repeat count (number of times access was attempted, as applicable). |
| Threat Prevention | Displays threat information grouped by the following attributes:<br>• Threats<br>• Types<br>• Spyware<br>• Spyware Phone Home<br>• Spyware Downloads<br>• Vulnerability<br>• Virus<br>Each chart can include the threat ID, count (number of occurrences), number of sessions, and subtype (such as vulnerability), as applicable |

 © Palo Alto Networks, Inc.

| ACC Chart | Description |
|---|---|
| Data Filtering | Displays information on data filtered by the firewall grouped by the following attributes:<br>• Content/File Types<br>• Types<br>• File Names |
| HIP Matches | Displays the host information collected by the firewall grouped by:<br>• HIP Objects<br>• HIP Profiles |

## ACC Detail Pages

To view additional details, click any of the links on the ACC charts. A details page opens to show information about the item at the top and additional lists for related items. For example, click on the web-browsing link on the Application chart opens the Application Information page for web-browsing:

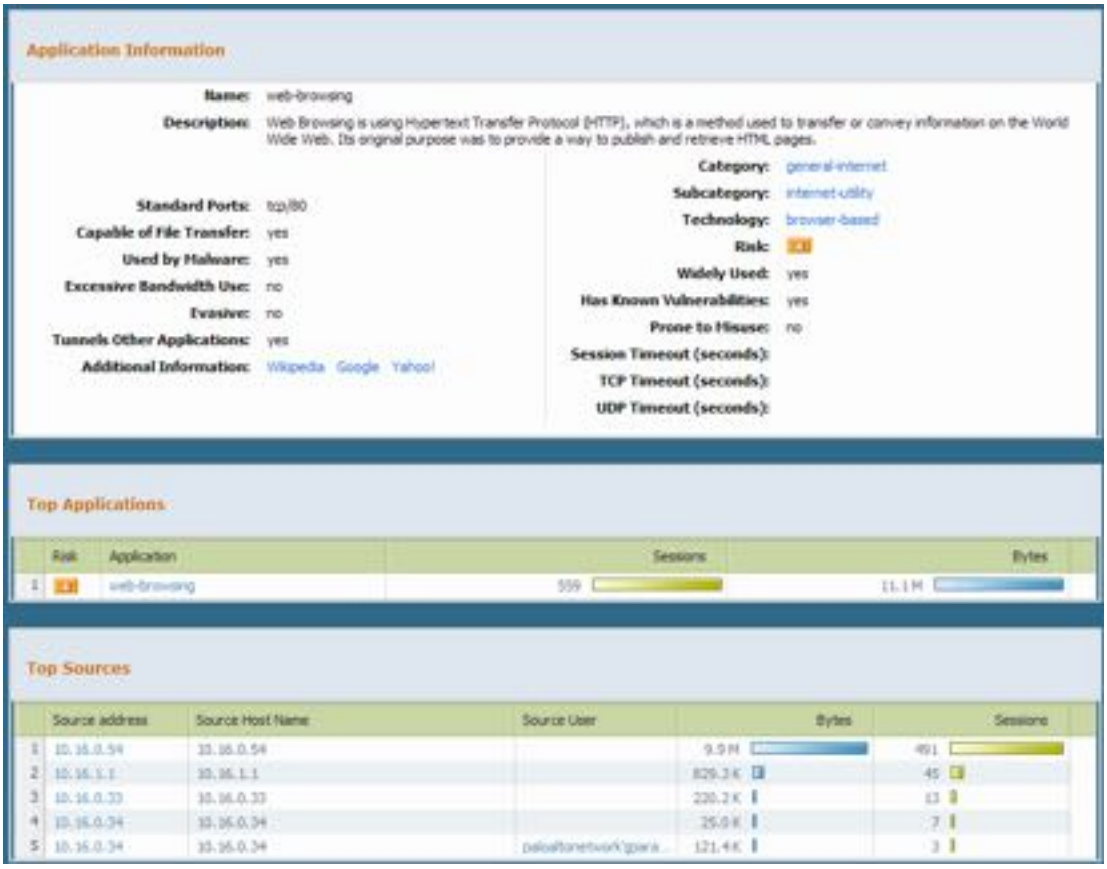


 © Palo Alto Networks, Inc.

## Use the ACC

The following procedure describes how to use the **ACC** tab and how to customize your view:

| Use the ACC | |
|---|---|
| Step 1 | On the **ACC**, change one or more of the settings at the top of the page.<br>• Use the drop-down to select Applications, URL Categories, Threats, Content/File Types, and HIP Objects to view.<br>• Select a virtual system, if virtual systems are defined.<br>• Select a time period from the **Time** drop-down. The default is Last Hour.<br>• Select a sorting method from the **Sort By** drop-down. You can sort the charts in descending order by number of sessions, bytes, or threats. The default is by number of sessions.<br>• For the selected sorting method, select the top number of applications and application categories shown in each chart from the **Top** drop-down.<br>Click the submit icon to apply the selected settings. |
| Step 2 | To open log pages associated with the information on the page, use the log links in the upper-right corner of the page, as shown here. The context for the logs matches the information on the page.<br> |
| Step 3 | To filter the list, click an item in one of the columns, this will add that item to the filter bar located above the log column names. After adding the desired filters, click the Apply Filter icon .<br> 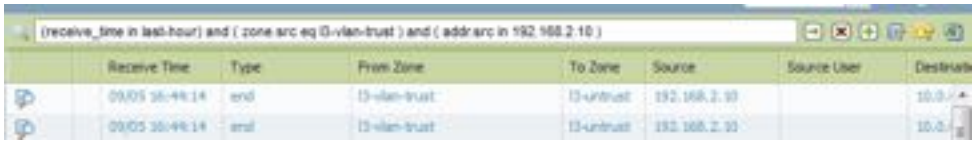 |

# Use App-Scope

The App-Scope reports provide visibility and analysis tools to help pinpoint problematic behavior, helping you understand changes in application usage and user activity, users and applications that take up most of the network bandwidth, and identify network threats.

With the App-Scope reports, you can quickly see if any behavior is unusual or unexpected. Each report provides a dynamic, user-customizable window into the network; hovering the mouse over and clicking either the lines or bars on the charts opens detailed information about the specific application, application category, user, or source on the ACC.

The following App-Scope reports are available:

- Summary Report
- Change Monitor Report
- Threat Monitor Report
- Threat Map Report
- Network Monitor Report
- Traffic Map Report


© Palo Alto Networks, Inc.

## Summary Report

The App-Scope Summary report displays charts for the top five gainers, losers, and bandwidth consuming applications, application categories, users, and sources.

## Change Monitor Report

The App-Scope Change Monitor report displays changes over a specified time period. For example, the following chart displays the top applications that gained in use over the last hour as compared with the last 24-hour period. The top applications are determined by session count and sorted by percent.

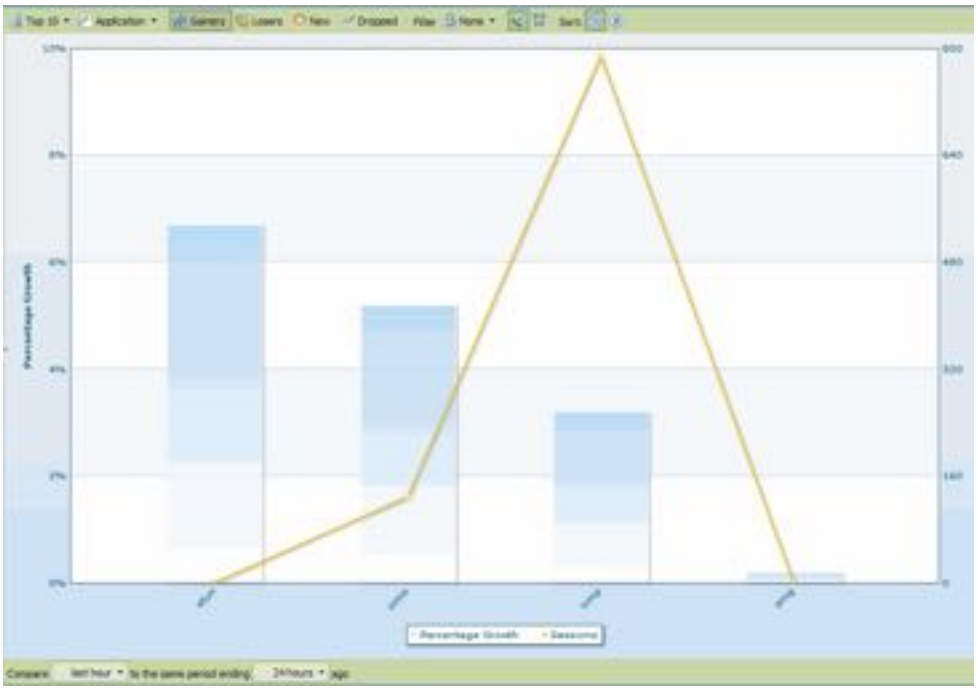

The Change Monitor Report contains the following buttons and options.

| Button | Description |
|---|---|
| Top 10 | Determines the number of records with the highest measurement included in the chart. |
| Application | Determines the type of item reported: Application, Application Category, Source, or Destination. |
| Gainers | Displays measurements of items that have increased over the measured period. |
| Losers | Displays measurements of items that have decreased over the measured period. |
| New | Displays measurements of items that were added over the measured period. |
| Dropped | Displays measurements of items that were discontinued over the measured period. |
| Filter None | Applies a filter to display only the selected item. **None** displays all entries. |


© Palo Alto Networks, Inc.

| Button | Description |
|---|---|
| | Determines whether to display session or byte information. |
| 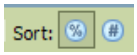 | Determines whether to sort entries by percentage or raw growth. |
| 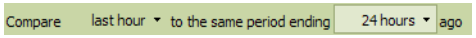 | Specifies the period over which the change measurements are taken. |

## Threat Monitor Report

The App-Scope Threat Monitor report displays a count of the top threats over the selected time period. For example, the following figure shows the top 10 threat types over the last 6 hours.

Each threat type is color-coded as indicated in the legend below the chart. The Threat Monitor report contains the following buttons and options.

| Button | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| | Determines the type of item measured: Threat, Threat Category, Source, or Destination. |
| | Applies a filter to display only the selected type of items. |
| | Determines whether the information is presented in a stacked column chart or a stacked area chart. |
| 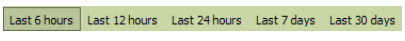 | Specifies the period over which the measurements are taken. |

 © Palo Alto Networks, Inc.

## Threat Map Report

The App-Scope Threat Map report shows a geographical view of threats, including severity. Each threat type is color-coded as indicated in the legend below the chart.

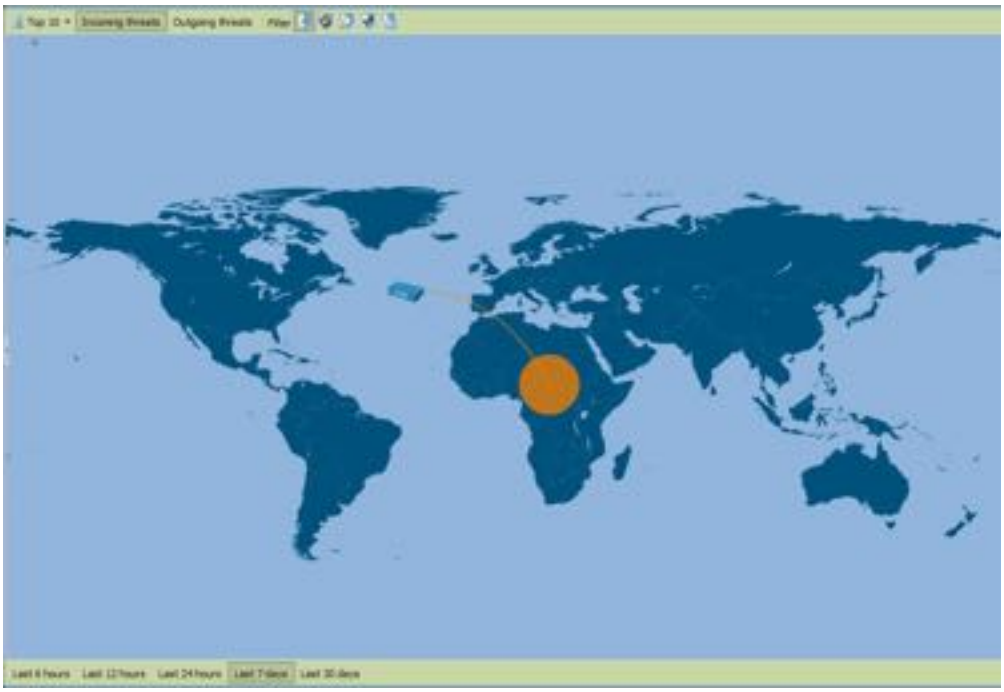

Click a country on the map to zoom in. Click the **Zoom Out** button in the lower right corner of the screen to zoom out. The Threat Map report contains the following buttons and options.

| Button | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| 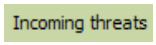  | Displays incoming threats. |
| 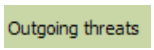  | Displays outgoing threats. |
| | Applies a filter to display only the selected type of items. |
| | Indicates the period over which the measurements are taken. |

## Network Monitor Report

The App-Scope Network Monitor report displays the bandwidth dedicated to different network functions over the specified period of time. Each network function is color-coded as indicated in the legend below the chart. For example, the image below shows application bandwidth for the past 7 days based on session information.

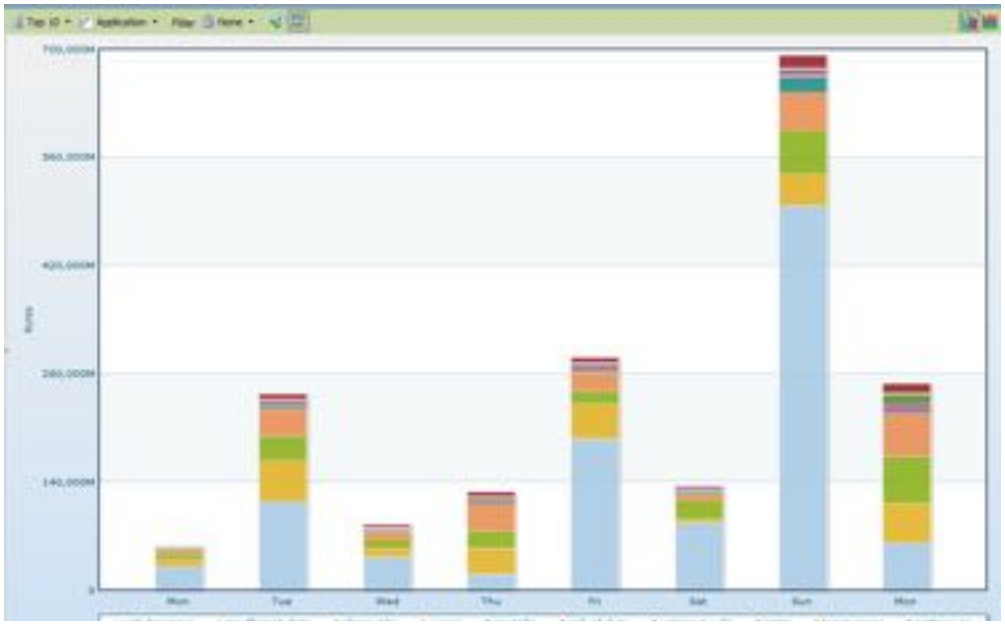

The Network Monitor report contains the following buttons and options.

| Button | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| 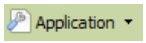 | Determines the type of item reported: Application, Application Category, Source, or Destination. |
| 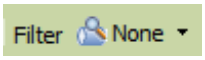 | Applies a filter to display only the selected item. **None** displays all entries. |
| | Determines whether to display session or byte information. |
| | Determines whether the information is presented in a stacked column chart or a stacked area chart. |
| 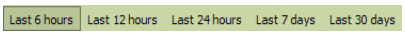 | Indicates the period over which the change measurements are taken. |

 © Palo Alto Networks, Inc.

## Traffic Map Report

The App-Scope Traffic Map report shows a geographical view of traffic flows according to sessions or flows.

Each traffic type is color-coded as indicated in the legend below the chart. The Traffic Map report contains the following buttons and options.

| Buttons | Description |
|---|---|
| Top 10 | Determines the number of records with the highest measurement included in the chart. |
| Incoming threats | Displays incoming threats. |
| Outgoing threats | Displays outgoing threats. |
| 010 101 | Determines whether to display session or byte information. |
| Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days | Indicates the period over which the change measurements are taken. |

# Take Packet Captures

PAN-OS supports packet capture for troubleshooting or detecting unknown applications. You can define filters such that only the packets that match the filters are captured. The packet captures are locally stored on the device and are available for download to your local computer.

Packet Capture is for troubleshooting only. This feature can cause the system performance to degrade and should be used only when necessary. Remember to disable the feature after you complete the packet capture.

The following table describes the packet capture settings on **Monitor > Packet Capture**.

| Field | Description |
|---|---|
| **Configure Filtering** | |
| **Manage Filters** | Click **Manage Filters**, click **Add** to add a new filter, and specify the following information:<br>• **Id**—Enter or select an identifier for the filter.<br>• **Ingress Interface**—Select the firewall interface.<br>• **Source**—Specify the source IP address.<br>• **Destination**—Specify the destination IP address.<br>• **Src Port**—Specify the source port.<br>• **Dest Port**—Specify the destination port.<br>• **Proto**—Specify the protocol to filter.<br>• **Non-IP**—Choose how to treat non-IP traffic (exclude all IP traffic, include all IP traffic, include only IP traffic, or do not include an IP filter).<br>• **IPv6**—Select the check box to include IPv6 packets in the filter. |
| **Filtering** | Click to toggle the filtering selections on or off. |
| **Pre-Parse Match** | Click to toggle the pre-parse match option on or off.<br>The pre-parse-match option is added for advanced troubleshooting purposes. After a packet enters the ingress port, it proceeds through several processing steps before it is parsed for matches against pre-configured filters.<br>It is possible for a packet, due to a failure, to not reach the filtering stage. This can occur, for example, if a route lookup fails.<br>Set the pre-parse-match setting to **ON** to emulate a positive match for every packet entering the system. This allows the firewall to capture even the packets that do not reach the filtering process. If a packet is able to reach the filtering stage, it is then processed according to the filter configuration and discarded if it fails to meet filtering criteria. |

 © Palo Alto Networks, Inc.

| Field | Description |
|---|---|
| **Configure Capturing** | |
| **Packet Capture** | Click to toggle packet capturing on or off.<br>For anti-spyware and vulnerability protection profiles, you can enable extended packet captures for rules and exceptions defined in the profile. This functionality allows the firewall to capture from 1 to 50 packets and provides more context when analyzing the threat logs.<br>To define the extended packet capture length:<br>1. Select **Device > Setup > Content-ID.**<br>2. Edit the Threat Detection Settings section to specify the **Capture Length** for the number of packets to capture.<br> 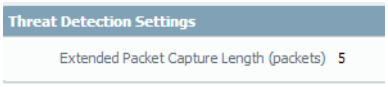 <br>3. View the packet capture in **Monitor > Logs > Threat**.<br>Locate the threat log entry and click the green arrow (Packet Capture) icon in the corresponding row to view the capture. |
| **Packet Capture Stage** | Select **Add** and specify the following:<br>• **Stage**—Indicate the point at which to capture the packet:<br>• **drop**—When packet processing encounters an error and the packet is to be dropped.<br>• **firewall**—When the packet has a session match or a first packet with a session is successfully created.<br>• **receive**—When the packet is received on the dataplane processor.<br>• **transmit**—When the packet is to be transmitted on the dataplane processor.<br>• **File**—Specify the capture file name. The file name should begin with a letter and can include letters, digits, periods, underscores, or hyphens.<br>• **Packet Count**—Specify the number of packets after which capturing stops.<br>• **Byte Count**—Specify the number of bytes after which capturing stops. |
| **Captured Files** | |
| **Captured Files** | Select **Delete** to remove a packet capture file from the list displaying captured files. |
| **Settings** | |
| **Clear All Settings** | Select **Clear All Settings** to clear all packet capture settings. |

# Monitor the Firewall

The following sections describe the methods you can use to monitor the firewall and provide basic setup instructions:

- Monitor Applications and Threats
- Monitor Log Data
- Monitor the Dashboard
- View Reports

You can also configure the firewall (excluding PA-4000 Series and PA-7050 firewalls) to export flow data to a NetFlow collector for analysis and reporting. To configure NetFlow Settings, refer to the PAN-OS-6.0 Web Interface Reference Guide.


© Palo Alto Networks, Inc.

## Monitor Applications and Threats

All Palo Alto Networks next-generation firewalls come equipped with the App-ID technology, which identifies the applications traversing your network, irrespective of protocol, encryption, or evasive tactic. You can then Use the Application Command Center to monitor the applications. ACC graphically summarizes the log database to highlight the applications traversing your network, who is using them, and their potential security impact. ACC is dynamically updated, using the continuous traffic classification that App-ID performs; if an application changes ports or behavior, App-ID continues to see the traffic, displaying the results in ACC.

You can quickly investigate new, risky, or unfamiliar applications that appear in ACC with a single click that displays a description of the application, its key features, its behavioral characteristics, and who is using it. Additional visibility into URL categories, threats, and data provides a complete and well-rounded picture of network activity. With ACC, you can very quickly learn more about the traffic traversing the network and then translate that information into a more informed security policy.

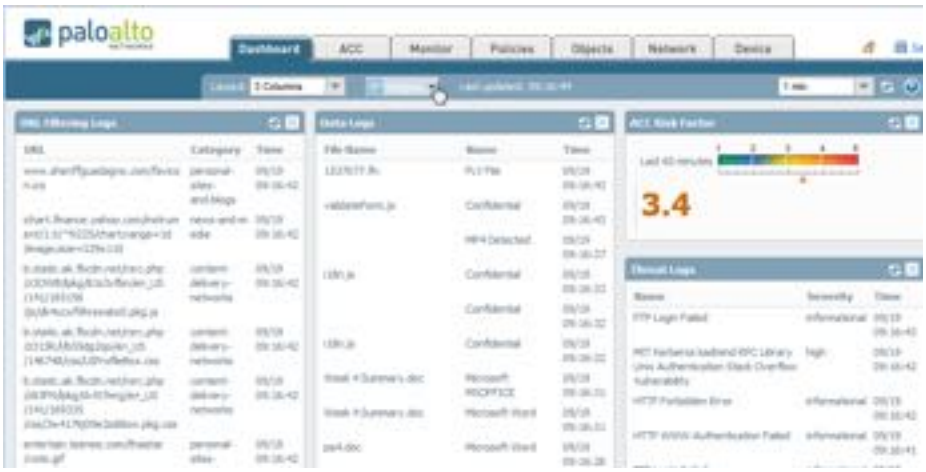

# Monitor Log Data

All Palo Alto Networks next-generation firewalls can generate log files that provide an audit trail of the activities and events on the firewall. There are separate logs for separate types of activities and events. For example, the Threat logs record all traffic that causes the firewall to generate a security alarm, whereas URL Filtering logs record all traffic that matches a URL Filtering profile attached to a security policy, and Config logs record all changes to the firewall configuration.

You can either Forward Logs to External Services or you can view logs locally on the device as follows:

- View the Log Files
- Filter Log Data

## View the Log Files

The firewall maintains logs for WildFire, configurations, system, alarms, traffic flows, threats, URL filtering, data filtering, and Host Information Profile (HIP) matches. You can view the current logs at any time. To locate specific entries, you can apply filters to most of the log fields.

The firewall displays the information in logs so that role-based administration permissions are respected. When you display logs, only the information that you have permission to see is included. For information on administrator permissions, see Administrative Roles.

By default all log files are generated and stored locally on the firewall. You can view these log files directly (**Monitor > Logs**):

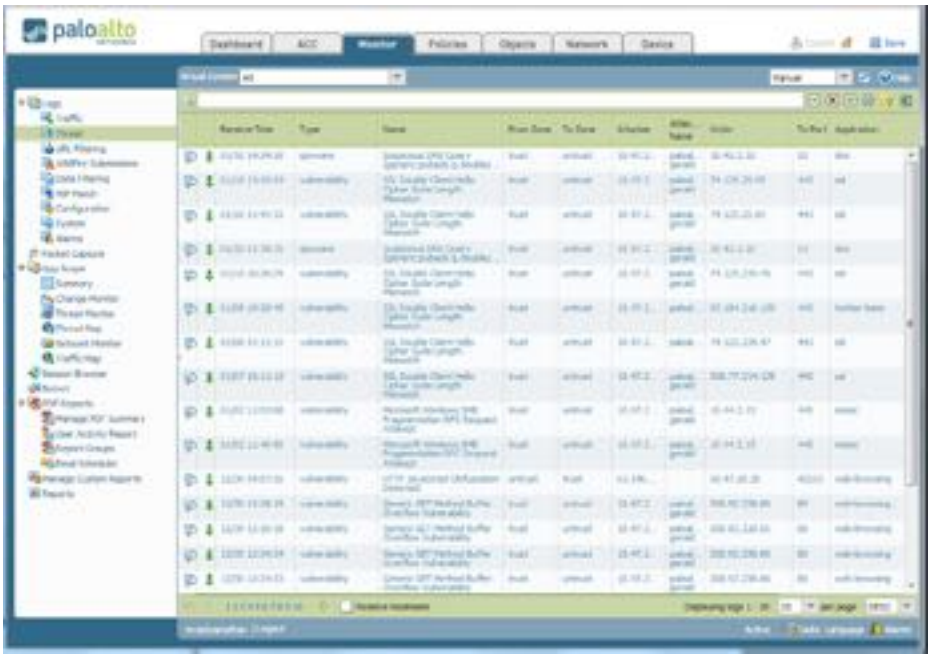

To display additional details, click the spyglass icon for an entry.

 © Palo Alto Networks, Inc.

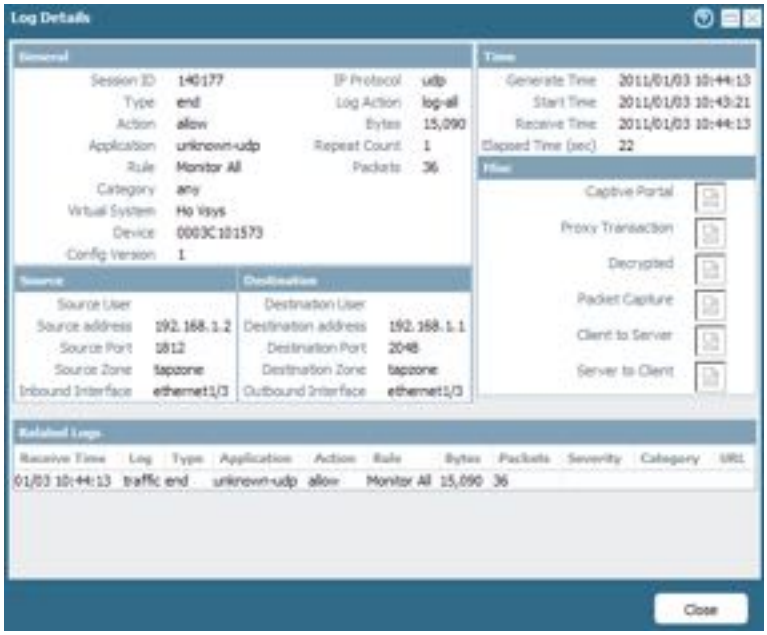


The following table includes information on each log type:

| Log Description Charts | Description |
|---|---|
| Traffic | Displays an entry for the start and end of each session. Each entry includes the date and time, the source and destination zones, addresses, and ports, the application name, the security rule name applied to the flow, the rule action (**allow, deny, or drop**), the ingress and egress interface, and the number of bytes.<br>Click next to an entry to view additional details about the session, such as whether an ICMP entry aggregates multiple sessions between the same source and destination (the **Count** value will be greater than one).<br>The **Type** column indicates whether the entry is for the start or end of the session, or whether the session was denied or dropped. A **drop** indicates that the security rule that blocked the traffic specified **any** application, while a **deny** indicates the rule identified a specific application.<br>If traffic is dropped before the application is identified, such as when a rule drops all traffic for a specific service, the application is shown as **not-applicable**. |

| Log Description Charts | Description |
|---|---|
| Threat | Displays an entry when traffic matches a Security Profile (Antivirus, Anti-spyware, Vulnerability, URL Filtering, File Blocking, Data Filtering, or DoS Protection) that is attached to a security policy on the firewall. Each entry includes the date and time, a threat name or URL, the source and destination zones, addresses, and ports, the application name, and the alarm action (**allow** or **block**) and severity.<br>Click next to an entry to view additional details about the threat, such as whether the entry aggregates multiple threats of the same type between the same source and destination (the **Count** value will be greater than one).<br>The **Type** column indicates the type of threat, such as "virus" or "spyware." The **Name** column is the threat description or URL, and the **Category** column is the threat category (such as "keylogger") or URL category.<br>If local packet captures are enabled, click next to an entry to access the captured packets. To enable local packet captures, see Take Packet Captures. |
| URL Filtering | Displays logs for all traffic that matches a URL Filtering profile attached to a security policy. For example, if policy blocks access to specific web sites and web site categories or if policy is configured to generate an alert when a web site is accessed. For information on defining URL filtering profiles, see URL Filtering. |
| WildFire Submissions | Displays logs for files that are uploaded and analyzed by the WildFire cloud, log data is sent back to the device after analysis, along with the analysis results. |
| Data Filtering | Displays logs for the security policies that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall. See Set Up Data Filtering for information on defining data filtering profiles.<br>This log also shows information for file blocking profiles. For example, if you are blocking *.exe* files, the log will show the files that were blocked. If you forward files to WildFire, you will see the results of that action. In this case, if you are forwarding PE files to WildFire, for example, the log will show that the file was forwarded and will also show the status on whether or not it was uploaded to WildFire successfully or not. |
| Configuration | Displays an entry for each configuration change. Each entry includes the date and time, the administrator username, the IP address from where the change was made, the type of client (XML, Web or CLI), the type of command executed, whether the command succeeded or failed, the configuration path, and the values before and after the change. |
| System | Displays an entry for each system event. Each entry includes the date and time, the event severity, and an event description. |
| HIP Match | Displays traffic flows that match a HIP Object or HIP Profile that you have configured. |

## Filter Log Data

Each log page has a filter area at the top of the page.

Use the filter area as follows:

 © Palo Alto Networks, Inc.

- Click any of the underlined links in the log listing to add that item as a log filter option. For example, if you click the **Host** link in the log entry for 10.0.0.252 and **Web Browsing** in both items are added, and the search will find entries that match both (AND search).
- To define other search criteria, click **Add Log Filter**. Select the type of search (and/or), the attribute to include in the search, the matching operator, and the values for the match, if appropriate. Click **Add** to add the criterion to the filter area on the log page, and then click **Close** to close the pop-up window. Click **Apply Filter** to display the filtered list.

If the **Value** string matches an **Operator** (such as **has** or **in**), enclose the string in quotation marks to avoid a syntax error. For example, if you filter by destination country and use `IN` as a **Value** to specify INDIA, enter the filter as `( dstloc eq "IN" )`.

You can combine filter expressions added on the log page with those you define in the Add Log Filter dialog. The filter field on the log page displays each filter as an entry.

If you add a **Receive Time** filter with the **Operator** set to **in** and the **Value** set to **Last 60 seconds**, some of the page links on the log viewer might not show results because the number of pages might grow or shrink due to the dynamic nature of the selected time.

- To clear filters and redisplay the unfiltered list, click **Clear Filter**.
- To save your selections as a new filter, click **Save Filter**, enter a name for the filter, and click **OK**.
- To export the current log listing (as shown on the page, including any applied filters) click **Save Filter**. Select whether to open the file or save it to disk, and select the check box if you want to always use the same option. Click **OK**.
- To export the current log listing in CSV Format, select the **Export to CSV** icon. By default, exporting the log listing to CSV format generates a CSV report with up to 2,000 rows of logs. To change the limit for rows displayed in CSV reports, use the **Max Rows in CSV Export** field on the **Log Export and Reporting** subtab (select **Device > Setup > Management > Logging and Reporting Settings**).

To change the automatic refresh interval, select an interval from the drop-down list (**1 min**, **30 seconds**, **10 seconds**, or **Manual**).

To change the number of log entries per page, select the number of rows from the **Rows** drop-down.

Log entries are retrieved in blocks of 10 pages. Use the paging controls at the bottom of the page to navigate through the log list. Select the **Resolve Hostname** check box to begin resolving external IP addresses to domain names.

## Monitor the Dashboard

You can also monitor the local log data directly from the **Dashboard** by adding the associated widgets:

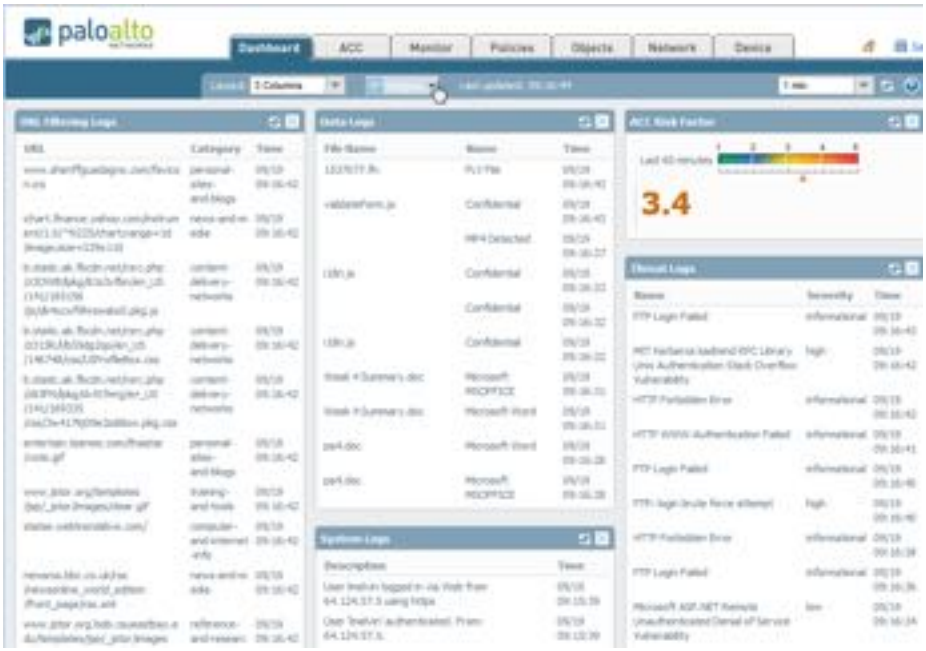


 © Palo Alto Networks, Inc.

## View Reports

The firewall also uses the log data to generate reports (**Monitor > Reports**) that display the log data in a tabular or graphical format. See About Reports for more details on the predefined and custom reports available on the firewall.

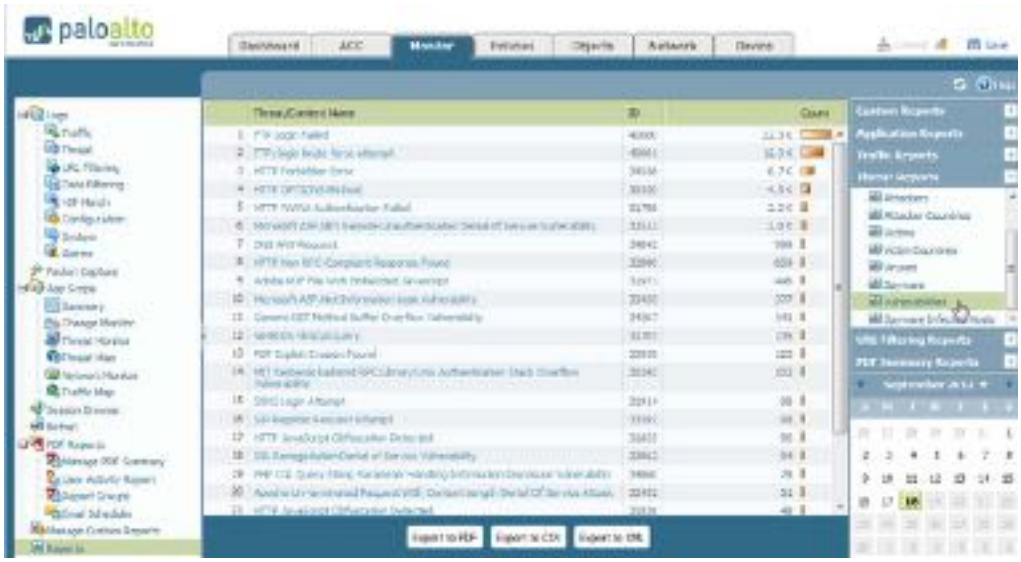

# Forward Logs to External Services

Depending on the type and severity of the data in the log files, you may want to be alerted to critical events that require your attention, or you may have policies that require you to archive the data for longer than it can be stored on the firewall. In these cases you will want to forward your log data to an external service for archive, notification, and/or analysis.

To forward log data to an external service you must complete the following tasks:

- Configure the firewall to access the remote services that will be receiving the logs. See Define Remote Logging Destinations.
- Configure each log type for forwarding. See Enable Log Forwarding.

 © Palo Alto Networks, Inc.

## Define Remote Logging Destinations

In order to reach an external service—such as a Syslog server or SNMP trap manager—the firewall must know the details of how to access and, if necessary, authenticate to the service. On the firewall, you define this information in a Server Profile. You must create a Server Profile for each external service you want the firewall to interact with. The type of logging destinations you need to set up and which logs you forward will depend on your needs. Some common log forwarding scenarios include the following:

- For immediate notification about critical system events or threats that require your attention, you can generate SNMP traps or send email alerts. See Set Up Email Alerts and/or Set Up SNMP Trap Destinations.
- For long-term storage and archival of data and for centralized device monitoring, you can send the log data to a Syslog server. See Define Syslog Servers. This enables integration with third-party security monitoring tools, such as Splunk! or ArcSight. You can also secure the channel between the firewall and the Syslog server. See Configure the Firewall to Authenticate to the Syslog Server.
- For aggregation and reporting of log data from multiple Palo Alto Networks firewalls, you can forward logs to a Panorama Manager or Panorama Log Collector. See Enable Log Forwarding.

You can define as many Server Profiles as you need. For example, you could use separate Server Profiles to send traffic logs to one Syslog server and system logs to a different one. Or, you could include multiple server entries in a single Server Profile to enable you to log to multiple Syslog servers for redundancy.

By default, all log data is forwarded over the MGT interface. If you plan to use an interface other than MGT, you will need to configure a Service Route for each service to which you plan to forward logs as described in Step 5 in the procedure to Set Up Network Access for External Services.

## Set Up Email Alerts

| Set Up Email Alerts | | |
|---|---|---|
| Step 1 | Create a Server Profile for your email server. | 1. Select **Device > Server Profiles > Email**.<br>2. Click **Add** and then enter a **Name** for the profile.<br>3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.<br>4. Click **Add** to add a new email server entry and enter the information required to connect to the Simple Mail Transport Protocol (SMTP) server and send email (you can add up to four email servers to the profile):<br>• **Server**—Name to identify the mail server (1-31 characters). This field is just a label and does not have to be the host name of an existing SMTP server.<br>• **Display Name**—The name to show in the From field of the email.<br>• **From**—The email address where notification emails will be sent from.<br>• **To**—The email address to which notification emails will be sent.<br>• **Additional Recipient**—If you want the notifications sent to a second account, enter the additional address here. You can only add one additional recipient. To add multiple recipients, add the email address of a distribution list.<br>• **Gateway**—The IP address or host name of the SMTP gateway to use to send the emails.<br>5. Click **OK** to save the server profile. |
| Step 2 | (Optional) Customize the format of the email messages the firewall sends. | Select the **Custom Log Format** tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide. |
| Step 3 | Save the server profile and commit your changes. | 1. Click **OK** to save the profile.<br>2. Click **Commit** to save the changes to the running configuration. |

## Set Up SNMP Trap Destinations

Simple Network Management Protocol (SNMP) is a standard facility for monitoring the devices on your network. You can configure the firewall to send SNMP traps to your SNMP management software to alert you to critical system events or threats that require your immediate attention.

You can also use SNMP to monitor the firewall. In this case, your SNMP manager must be configured to get statistics from the firewall rather than (or in addition to) having the firewall send traps to the manager. For more information, see Configure the Firewall to Authenticate to the Syslog Server.

 © Palo Alto Networks, Inc.

**Set Up SNMP Trap Destinations**

Step 1 (SNMP v3 only) Get the engine ID for the firewall.

In many cases, the MIB browser or SNMP manager will automatically discover the engine ID upon successful connection to the SNMP agent on the firewall. You can usually find this information in the agent settings section of the interface. Refer to the documentation for your specific product for instructions on finding the agent information.

In order to find out the firewall's engine ID, you must configure the firewall for SNMP v3 and send a GET message from your SNMP manager or MIB browser as follows:

1. Enable the interface to allow inbound SNMP requests:
   - If you will be receiving SNMP GET messages on the MGT interface, select **Device > Setup > Management** and click the Edit icon in the Management Interface Settings section of the screen. In the **Services** section, select the **SNMP** check box and then click **OK**.
   - If you will be receiving SNMP GET messages on a different interface, you must associate a management profile with the interface and enable SNMP management.
2. Configure the firewall for SNMP v3 as described in Step 2 in Set Up SNMP Monitoring. If you do not configure the firewall for SNMP v3 your MIB browser will not allow you to GET the engine ID.
3. Connect your MIB browser or SNMP manager to the firewall and run a GET for OID 1.3.6.1.6.3.10.2.1.1.0. The value that is returned is the unique engine ID for the firewall.

Step 2 Create a Server Profile that contains the information for connecting and authenticating to the SNMP manager(s).

1. Select **Device > Server Profiles > SNMP Trap**.
2. Click **Add** and then enter a **Name** for the profile.
3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.
4. Specify the version of SNMP you are using (V2c or V3).
5. Click **Add** to add a new **SNMP Trap Receiver** entry (you can add up to four trap receivers per server profile). The required values depend on whether you are using SNMP V2c or V3 as follows:

   SNMP V2c

   - **Server**—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the host name of an existing SNMP server.
   - **Manager**—The IP address of the SNMP manager to which you want to send traps.
   - **Community**—The community string required to authenticate to the SNMP manager.

   SNMP V3

   - **Server**—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the host name of an existing SNMP server.
   - **Manager**—The IP address of the SNMP manager to which you want to sent traps.
   - **User**—The username required to authenticate to the SNMP manager.
   - **EngineID**—The engine ID of the firewall, as identified in Step 1. This is a hexadecimal value from 5 to 64 bytes with a 0x prefix. Each firewall has a unique engine ID.
   - **Auth Password**—The password to be used for authNoPriv level messages to the SNMP manager. This password will be hashed using Secure Hash Algorithm (SHA-1), but will not be encrypted.
   - **Priv Password**—The password to be used for authPriv level messages to the SNMP manager. This password be hashed using SHA and will be encrypted using Advanced Encryption Standard (AES 128).
6. Click **OK** to save the server profile.

| Set Up SNMP Trap Destinations (Continued) | | |
|---|---|---|
| Step 3 | (Optional) Set up a service route for SNMP traps. | By default, SNMP traps are sent over the MGT interface. If you want to use a different interface for SNMP traps, you must edit the service route to enable the firewall to reach your SNMP manager. See Set Up Network Access for External Services for instructions. |
| Step 4 | Commit your changes. | Click **Commit**. The device may take up to 90 seconds to save your changes. |
| Step 5 | Enable the SNMP manager to interpret the traps it receives from the firewall. | Load the PAN-OS MIB files into your SNMP management software and compile them. Refer to the documentation for your SNMP manager for specific instructions on how to do this. |

## Define Syslog Servers

Syslog is a standard log transport mechanism that enables the aggregation of log data from different network devices—such as routers, firewalls, printers—from different vendors into a central repository for archive, analysis, and reporting.

The firewall generates five types of logs that can be forwarded to an external syslog server: traffic, threat, WildFire, host information profile (HIP) match, config, and system. If you want to forward all or some of these logs to an external service for log-term storage and analysis, you can use TCP or SSL for reliable and secure transport of logs, or UDP for non-secure transport.

 © Palo Alto Networks, Inc.

| Set Up Syslog Forwarding | |
|---|---|
| Step 1 Create a Server Profile that contains the information for connecting to the Syslog server(s). | 1. Select **Device > Server Profiles > Syslog**.<br>2. Click **Add** and then enter a **Name** for the profile.<br>3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.<br>4. Click **Add** to add a new Syslog server entry and enter the information required to connect to the Syslog server (you can add up to four Syslog servers to the same profile):<br>• **Name**—Unique name for the server profile.<br>• **Server**—IP address or fully qualified domain name (FQDN) of the Syslog server.<br>• **Transport**—Select TCP, UDP, or SSL as the method of communication with the syslog server.<br>• **Port**—The port number on which to send Syslog messages (default is UDP on port 514); you must use the same port number on the firewall and the Syslog server.<br>• **Format**—Select the Syslog message format to use, BSD or IETF. Traditionally, BSD format is over UDP and IETF format is over TCP/SSL. For setting up secure syslog forwarding with client authentication, see Configure the Firewall to Authenticate to the Syslog Server.<br>• **Facility**—Select one of the Syslog standard values, which is used to calculate the priority (PRI) field in your Syslog server implementation. You should select the value that maps to how you use the PRI field to manage your Syslog messages.<br>5. (Optional) To customize the format of the Syslog messages the firewall sends, select the **Custom Log Format** tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide.<br>6. Click **OK** to save the server profile. |

| Set Up Syslog Forwarding (Continued) | |
|---|---|
| Step 2 (Optional) Configure the header format used in Syslog messages. Choosing the header format offers more flexibility in filtering and reporting on the log data for some SIEMs.<br> 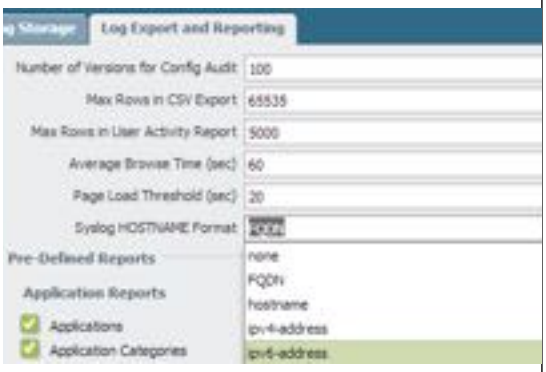 <br><br>This is a global setting and applies to all syslog server profiles configured on the appliance. | 1. Select **Device > Setup > Management** and click the Edit icon in the Logging and Reporting Settings section.<br>2. Select **Log Export and Reporting**.<br>3. Select one of the following options from the **Send Hostname in Syslog** drop-down:<br>• **FQDN**— (the default) Concatenates the hostname and domain name defined on the sending device.<br>• **hostname**— Uses the hostname defined on the sending device.<br>• **ipv4-address**—Uses theIPv4 address of the interface used to send logs on the device. By default, this is the MGT interface of the device.<br>• **ipv6-address**—Uses the IPv6 address of the interface used to send logs on the device. By default, this is the MGT interface of the device.<br>• **none**—Leaves the hostname field unconfigured on the device. There is no identifier for the device that sent the logs.<br>4. Click **OK** and **Commit**. |
| Step 3 Commit your changes. | Click **Commit**. The device may take up to 90 seconds to save your changes. |
| Step 4 Enable log forwarding. | See Enable Log Forwarding.<br>You must configure each log type for forwarding and specify the severity for which the event is logged.<br>WildFire logs are a type of threat log, but they are not logged and forwarded along with threat logs. While WildFire logs use the same syslog format as threat logs, the threat subtype is preset to WildFire. Therefore, you must enable logging/forwarding for WildFire logs distinctly from threat logs. |
| Step 5 Review the logs on the syslog server. | To parse the logs, see Syslog Field Descriptions. |

## Configure the Firewall to Authenticate to the Syslog Server

To enable client authentication for syslog over SSL, you can use a trusted CA or a self-signed CA for generating certificates that can be used for secure syslog communication. Check for the following when generating a certificate for secure syslog communication:

- The private key must be available on the sending device; the keys cannot be stored on a Hardware Security Module (HSM).
- The subject and the issuer for the certificate must not be identical.

 © Palo Alto Networks, Inc.

- The certificate is neither a trusted CA nor a certificate signing request (CSR). Neither of these types of certificates can be enabled for secure syslog communication.

**Configure the Firewall to Authenticate to the Syslog Server**

Step 1 If the syslog server requires client authentication, generate the certificate for secure communication. For details on certificates, see Certificate Management.

To verify that the sending device is authorized to communicate with the syslog server, you must enable the following:

- The server and the sending device must have certificates that are signed by the Enterprise CA; or you can generate a self-signed certificate on the firewall, export the CA root certificate from the firewall and import it in to the syslog server.
- Use the Enterprise CA or the self-signed certificate to generate a certificate with the IP address of the sending device (as the Common Name) and enabled for use in secure syslog communication. The syslog server uses this certificate to verify that the firewall is authorized to communicate with the syslog server.

  Use the following steps to generate the certificate on the firewall or Panorama:

1. Select **Device > Certificate Management > Certificates > Device Certificates**.
2. Click **Generate** to create a new certificate that will be signed by a trusted CA or the self-signed CA.
3. Enter a **Name** for the certificate.
4. In **Common Name**, enter the IP address of the device sending logs to the syslog server.
5. Select **Shared** if you want the certificate to be a shared certificate on Panorama or to be shared by all virtual systems in a multiple virtual system firewall.
6. In **Signed by**, select the trusted CA or the self-signed CA that is trusted by both the syslog server and the sending device.
7. Click **Generate**. The certificate and the keypair will be generated.
8. Click the link with name of the certificate and enable the option Certificate for Secure Syslog for secure access to the syslog server.
9. **Commit** the changes.
10. Verify the certificate details and that it is marked for **Usage** as **Certificate for Secure Syslog**.

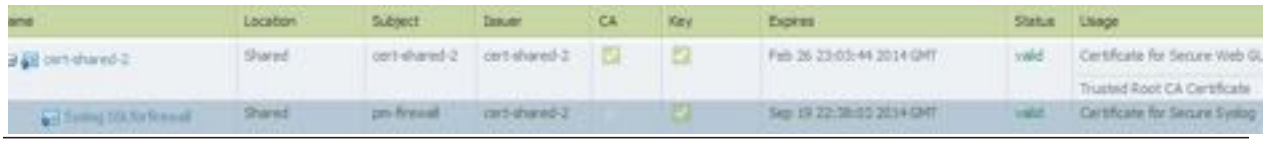

## Enable Log Forwarding

After you create the Server Profiles that define where to send your logs (see Define Remote Logging Destinations), you must enable log forwarding. For each log type, you can specify whether to forward it to Syslog, email, SNMP trap receiver, and/or Panorama.

Before you can forward log files to a Panorama Manager or a Panorama Log Collector, the firewall must be configured as a managed device. You can then enable log forwarding to Panorama for each type of log. For logs forwarded to Panorama, support for centralized log forwarding to an external syslog server is available.

You can use Secure Copy (SCP) commands from the CLI to export the entire log database (logdb) to an SCP server and import it to another firewall: refer to the PAN-OS Command Line Interface (CLI) Reference Guide. Because the log database is too large for an export or import to be practical on the following platforms, they do not support these options: PA-7050 firewall (all PAN-OS releases), Panorama virtual appliance running Panorama 6.0 or later releases, and Panorama M-Series appliances (all Panorama releases).

The way you enable forwarding depends on the log type:

- **Traffic Logs**—You enable forwarding of Traffic logs by creating a Log Forwarding Profile (**Objects > Log Forwarding**) and adding it to the security policies you want to trigger the log forwarding. Only traffic that matches a specific rule within the security policy will be logged and forwarded.
- **Threat Logs**—You enable forwarding of Threat logs by creating a Log Forwarding Profile (**Objects > Log Forwarding**) that specifies which severity levels you want to forward and then adding it to the security policies for which you want to trigger the log forwarding. A Threat log entry will only be created (and therefore forwarded) if the associated traffic matches a Security Profile (Antivirus, Anti-spyware, Vulnerability, URL Filtering, File Blocking, Data Filtering, or DoS Protection). The following table summarizes the threat severity levels:

| Severity | Description |
|---|---|
| **Critical** | Serious threats such as those that affect default installations of widely deployed software, result in root compromise of servers, and the exploit code is widely available to attackers. The attacker usually does not need any special authentication credentials or knowledge about the individual victims and the target does not need to be manipulated into performing any special functions. |
| **High** | Threats that have the ability to become critical but have mitigating factors; for example, they may be difficult to exploit, do not result in elevated privileges, or do not have a large victim pool. |
| **Medium** | Minor threats in which impact is minimized, such as DoS attacks that do not compromise the target or exploits that require an attacker to reside on the same LAN as the victim, affect only non-standard configurations or obscure applications, or provide very limited access. In addition, WildFire Submissions log entries with a malware verdict are logged as Medium. |
| **Low** | Warning-level threats that have very little impact on an organization's infrastructure. They usually require local or physical system access and may often result in victim privacy or DoS issues and information leakage. Data Filtering profile matches are logged as Low. |

 © Palo Alto Networks, Inc.

| Severity | Description |
|---|---|
| **Informational** | Suspicious events that do not pose an immediate threat, but that are reported to call attention to deeper problems that could possibly exist. URL Filtering log entries and WildFire Submissions log entries with a benign verdict are logged as Informational. |

- **Config Logs**—You enable forwarding of Config logs by specifying a Server Profile in the log settings configuration. (**Device > Log Settings > Config Logs**).
- **System Logs**—You enable forwarding of System logs by specifying a Server Profile in the log settings configuration. (**Device > Log Settings > System Logs**). You must select a Server Profile for each severity level you want to forward. For a partial list of system log messages and their corresponding severity levels, refer to the System Log Reference. The following table summarizes the system log severity levels:

| Severity | Description |
|---|---|
| **Critical** | Hardware failures, including HA failover and link failures. |
| **High** | Serious issues, including dropped connections with external devices, such as LDAP and RADIUS servers. |
| **Medium** | Mid-level notifications, such as antivirus package upgrades. |
| **Low** | Minor severity notifications, such as user password changes. |
| **Informational** | Log in/log off, administrator name or password change, any configuration change, and all other events not covered by the other severity levels. |

# Monitor the Firewall Using SNMP

All Palo Alto Networks firewalls support standard networking SNMP management information base (MIB) modules as well as proprietary Enterprise MIB modules. You can configure an SNMP manager to get statistics from the firewall. For example, you could configure your SNMP manager to monitor the interfaces, active sessions, concurrent sessions, session utilization percentage, temperature, and/or system uptime on the firewall.

Palo Alto Networks firewalls support SNMP GET requests only; SNMP SET requests are not supported.

| Set Up SNMP Monitoring | | |
|---|---|---|
| Step 1 | Enable the interface to allow inbound SNMP requests. | • If you will be receiving SNMP GET messages on the MGT interface, select **Device > Setup > Management** and click the Edit icon in the Management Interface Settings section of the screen. In the **Services** section, select the **SNMP** check box and then click **OK**.<br>• If you will be receiving SNMP GET messages on a different interface, you must associate a management profile with the interface and enable SNMP management. |
| Step 2 | From the web interface on the firewall, configure the settings to allow the SNMP agent on the firewall to respond to incoming GET requests from the SNMP manager. | 1. Select **Device > Setup > Operations > SNMP Setup**.<br>2. Specify the **Physical Location** of the firewall and the name or email address of an administrative **Contact**.<br>3. Select the SNMP **Version** and then enter the configuration details as follows (depending on which SNMP version you are using) and then click **OK**:<br>• **V2c**—Enter the **SNMP Community String** that will allow the SNMP manager access to the SNMP agent on the firewall. The default value is **public**, however because this is a well-known community string, it is a best practice to use a value that is not easily guessed.<br>• **V3**—You must create at least one View and one User in order to use SNMPv3. The view specifies which management information the manager has access to. If you want to allow access to all management information, just enter the top-level **OID** of .1.3.6.1 and specify the **Option** as **include** (you can also create views that exclude certain objects). Use **0xf0** as the **Mask**. Then when you create a user, select the **View** you just created and specify the **Auth Password** and **Priv Password**.<br>The authentication settings (the community string for V2c or the username and passwords for V3) configured on the firewall must match the value configured on the SNMP manager.<br>4. Click **OK** to save the settings.<br>5. Click **Commit** to save the SNMP settings. |
| Step 3 | Enable the SNMP manager to interpret firewall statistics. | Load the PAN-OS MIB files into your SNMP management software and, if necessary, compile them. Refer to the documentation for your SNMP manager for specific instructions on how to do this. |


© Palo Alto Networks, Inc.

| Set Up SNMP Monitoring (Continued) | | |
|---|---|---|
| Step 4 | Identify the statistics you want to monitor. | Using a MIB browser, walk the PAN-OS MIB files to identify the object identifiers (OIDs) that correspond to the statistics you want to monitor. For example, suppose you want to monitor Session Utilization Percentage on the firewall. Using a MIB browser you will see that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.1.2.3.1.0 in the PAN-COMMON-MIB. |
| Step 5 | Configure the SNMP management software to monitor the OIDs you are interested in. | Refer to the documentation for your SNMP manager for specific instructions on how to do this. |
| Step 6 | After you complete the configuration on both the firewall and the SNMP manager, you can begin monitoring the firewall from your SNMP management software. | The following is an example of how an SNMP manager displays real-time session utilization percentage statistics for a monitored PA-500 firewall:<br> 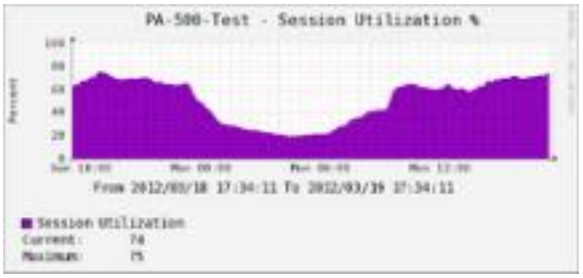 |

# Monitor the Firewall Using NetFlow

NetFlow is an industry-standard protocol that enables the firewall to record statistics on the IP traffic that traverses its interfaces. The firewall exports the statistics as NetFlow fields to a NetFlow collector. The NetFlow collector is a server you use to analyze network traffic for security, administration, accounting and troubleshooting purposes. All firewalls support NetFlow Version 9 except the PA-4000 Series and PA-7050 firewalls. The firewalls support only unidirectional NetFlow, not bidirectional. You can enable NetFlow exports on all interface types except HA, log card, or decrypt mirror. To identify firewall interfaces in a NetFlow collector, see Identify Firewall Interfaces in External Monitoring Systems.

The firewall supports standard and enterprise (PAN-OS specific) NetFlow templates. NetFlow collectors require templates to decipher the exported fields. The firewall selects a template based on the type of data it exports: IPv4 or IPv6 traffic, with or without NAT, and with standard or enterprise-specific fields. The firewall periodically refreshes templates to apply any changes to them. You configure the refresh frequency according to the requirements of the particular NetFlow collector you use.

| **Monitor the Firewall Using NetFlow** | | |
|---|---|---|
| Step 1 | Create a NetFlow server profile. | 1. Select **Device > Server Profiles > NetFlow** and click **Add**.<br>2. Enter a **Name** for the profile.<br>3. Specify the frequency at which the firewall refreshes templates (in **Minutes** or **Packets**) and exports records (**Active Timeout** in minutes).<br>4. Select the **PAN-OS Field Types** check box if you want the firewall to export App-ID and User-ID fields.<br>5. For each NetFlow collector that will receive fields (up to two per profile), click **Add** and enter an identifying server **Name**, hostname or IP address (**NetFlow Server**), and access **Port** (default 2055).<br>6. Click **OK** to save the profile. |
| Step 2 | Assign the NetFlow server profile to the interfaces that carry the traffic you want to analyze.<br>In this example, you assign the profile to an existing Layer 3 interface. | 1. Select **Network > Interfaces > Ethernet**.<br>2. Click an **Interface** name to edit it.<br>3. In the **NetFlow Profile** drop-down, select the NetFlow server profile.<br>4. Click **OK** and **Commit** your changes. |


© Palo Alto Networks, Inc.

# Identify Firewall Interfaces in External Monitoring Systems

When you use a NetFlow collector (see Monitor the Firewall Using NetFlow) or SNMP manager (see Monitor the Firewall Using SNMP) to monitor traffic flows, an interface index (SNMP ifindex object) identifies the firewall interface that carried a particular flow. The formula that the Palo Alto Networks firewall uses to calculate interface indexes varies by platform and whether the interface is physical or logical.

You cannot use SNMP to monitor logical interfaces, only physical interfaces. You can use NetFlow to monitor logical or physical interfaces.

Most NetFlow collectors use SNMP to determine the name of a physical interface based on the SNMP interface index.

Physical interface indexes have a range of 1-9999, which the firewall calculates as follows:

| Firewall Platform | Calculation | Example Interface Index |
|---|---|---|
| Non-chassis based:<br>VM-Series, PA-200, PA-500, PA-3000 Series, PA-4000 Series, PA-5000 Series<br> The PA-4000 Series platform supports SNMP but not NetFlow. | MGT port + physical port offset<br>• MGT port—This is a constant that depends on the platform:<br>  • 2 for hardware-based firewalls (for example, the PA-5000 Series firewall)<br>  • 1 for the VM-Series firewall<br>• Physical port offset—This is the physical port number. | PA-5000 Series firewall, Eth1/4 =<br>2 (MGT port) + 4 (physical port) = **6** |
| Chassis based:<br>PA-7050 firewalls<br> This platform supports SNMP but not NetFlow. | (Max. ports * slot) + physical port offset + MGT port<br>• Maximum ports—This is a constant of 64.<br>• Slot—This is the chassis slot number of the network interface card.<br>• Physical port offset—This is the physical port number.<br>• MGT port—This is a constant of 5 for PA-7050 firewalls. | PA-7050 firewall, Eth3/9 =<br>[64 (max. ports) * 3 (slot)] + 9 (phyical port) + 5 (MGT port) = **206** |

Logical interface indexes for all platforms are nine-digit numbers that the firewall calculates as follows:

| Interface Type | Range | Digit 9 | Digits 7-8 | Digits 5-6 | Digits 1-4 | Example Interface Index |
|---|---|---|---|---|---|---|
| Layer 3 subinterface | 101010001-199999999 | Type: 1 | Interface slot: 1-9 (01-09) | Interface port: 1-9 (01-09) | Subinterface: suffix 1-9999 (0001-9999) | Eth1/5.22 = 100000000 (type) + 100000 (slot) + 50000 (port) + 22 (suffix) = **101050022** |
| Layer 2 subinterface | 101010001-199999999 | Type: 1 | Interface slot: 1-9 (01-09) | Interface port: 1-9 (01-09) | Subinterface: suffix 1-9999 (0001-9999) | Eth2/3.6 = 100000000 (type) + 200000 (slot) + 30000 (port) + 6 (suffix) = **102030006** |

| Interface Type | Range | Digit 9 | Digits 7-8 | Digits 5-6 | Digits 1-4 | Example Interface Index |
|---|---|---|---|---|---|---|
| Vwire subinterface | 101010001-199999999 | Type: 1 | Interface slot: 1-9 (01-09) | Interface port: 1-9 (01-09) | Subinterface: suffix 1-9999 (0001-9999) | Eth4/2.312 = 100000000 (type) + 400000 (slot) + 20000 (port) + 312 (suffix) = **104020312** |
| VLAN | 200000001-200009999 | Type: 2 | 00 | 00 | VLAN suffix: 1-9999 (0001-9999) | VLAN.55 = 200000000 (type) + 55 (suffix) = **200000055** |
| Loopback | 300000001-300009999 | Type: 3 | 00 | 00 | Loopback suffix: 1-9999 (0001-9999) | Loopback.55 = 300000000 (type) + 55 (suffix) = **300000055** |
| Tunnel | 400000001-400009999 | Type: 4 | 00 | 00 | Tunnel suffix: 1-9999 (0001-9999) | Tunnel.55 = 400000000 (type) + 55 (suffix) = **400000055** |
| Aggregate group | 500010001-500089999 | Type: 5 | 00 | AE suffix: 1-8 (01-08) | Subinterface: suffix 1-9999 (0001-9999) | AE5.99 = 500000000 (type) + 50000 (AE Suffix) + 99 (suffix) = **500050099** |


© Palo Alto Networks, Inc.

# Manage Reporting

The reporting capabilities on the firewall allow you to keep a pulse on your network, validate your policies, and focus your efforts on maintaining network security for keeping your users safe and productive.


- About Reports
- View Reports
- Disable Predefined Reports
- Generate Custom Reports
- Generate Botnet Reports
- Manage PDF Summary Reports
- Generate User/Group Activity Reports
- Manage Report Groups
- Schedule Reports for Email Delivery


© Palo Alto Networks, Inc.

## About Reports

The firewall includes predefined reports that you can use as-is, build custom reports that meet your needs for specific data and actionable tasks, or combine predefined and custom reports to compile information you need. The firewall provides the following types of reports:

- **Predefined Reports—**Allow you to view a quick summary of the traffic on your network. A suite of predefined reports are available in four categories —Applications, Traffic, Threat, and URL Filtering. See View Reports.
- **User or Group Activity Reports—**Allow you to schedule or create an on-demand report on the application use and URL activity for a specific user or for a user group; the report includes the URL categories and an estimated browse time calculation for individual users. See Generate User/Group Activity Reports.
- **Custom Reports**—Create and schedule custom reports that show exactly the information you want to see by filtering on conditions and columns to include. You can also include query builders for more specific drill down on report data. See Generate Custom Reports.
- **PDF Summary Reports—**Aggregate up to 18 predefined or custom reports/graphs from Threat, Application, Trend, Traffic, URL Filtering categories into one PDF document. See Manage PDF Summary Reports.
- **Botnet Reports**—Allow you to use behavior-based mechanisms to identify potential botnet-infected hosts in the network. See Generate Botnet Reports.
- **Report Groups**—Combine custom and predefined reports into report groups and compile a single PDF that is emailed to one or more recipients. See Manage Report Groups.

Reports can be generated on demand, on a recurring schedule, and can be scheduled for email delivery.

 © Palo Alto Networks, Inc.

# View Reports

The firewall provides an assortment of over 40 predefined reports that are generated everyday; these reports can be viewed directly on the firewall. In addition to these reports, you can view *scheduled* custom reports and summary reports.

About 200 MB of storage is allocated for saving reports on the firewall. This storage space is not user configurable, and older reports are purged to store recent reports. Therefore, for long term retention of reports, you can either export the reports or schedule the reports for email delivery. To disable selected reports and conserve system resources on the firewall, see Disable Predefined Reports.

User/group activity reports must be generated on demand or scheduled for email delivery. Unlike the other reports, these reports cannot be saved on the firewall.

**View Reports**

Step 1 Select **Monitor > Reports**.

The reports are chunked into sections on the right-hand side of the window: **Custom Reports**, **Application Reports**, **Traffic Reports**, **Threat Reports**, **URL Filtering Reports**, and **PDF Summary Reports**.

Step 2 Select a report to view. When you select a report, the previous day's report is displayed on screen.

To view reports for any of the previous days, select an available date from the calendar at the bottom of the page and select a report within the same section. If you change sections, the time selection is reset.

Step 3 To view a report offline, you can export the report to PDF, CSV or to XML formats. Click **Export to PDF**, **Export to CSV**, or **Export to XML** at the bottom of the page. Then print or save the file.

## Disable Predefined Reports

The firewall includes about 40 predefined reports that are automatically generated each day. If you do not use some or all of these predefined reports, you can disable selected reports and conserve system resources on the firewall.

Before disabling one or more predefined reports, make sure that the report is not included in a Group Report or a PDF Report. If the disabled predefined report is included in a group or PDF report, the Group/PDF report will be rendered without any data.

### Disable Predefined Reports

1. Select **Device > Setup > Management** on the firewall.
2. Click the Edit icon in the Logging and Reporting Settings section and select the **Log Export and Reporting** tab.
3. To disable reports:
   - Clear the check box corresponding to each report that you want to disable.
   - Select **Deselect All** to disable all predefined reports.

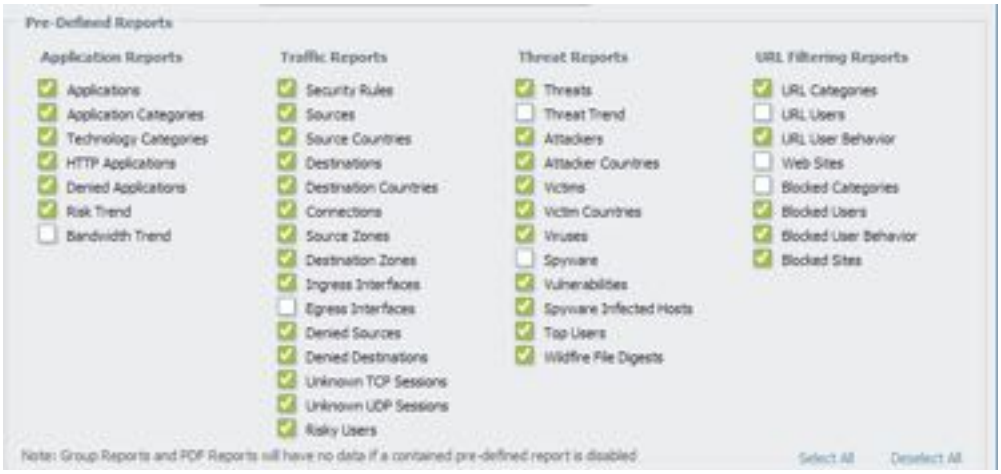


4. Click **OK**, and **Commit** the changes.

 © Palo Alto Networks, Inc.

## Generate Custom Reports

In order to create purposeful custom reports, you must consider the attributes or key pieces of information that you want to retrieve and analyze. This consideration guides you in making the following selections in a custom report:

| Selection | Description |
|---|---|
| Data Source | The data file that is used to generate the report. The firewall offers two types of data sources—Summary databases and Detailed logs.<br>• Summary databases are available for traffic, threat, and application statistics. The firewall aggregates the detailed logs on traffic, application, and threat at 15-minute intervals. The data is condensed—duplicate sessions are grouped together and incremented with a repeat counter, and some attributes (or columns) are not included in the summary—to allow faster response time when generating reports.<br>• Detailed logs are itemized and are a complete listing of all the attributes (or columns) that pertain to the log entry. Reports based on detailed logs take much longer to run and are not recommended unless absolutely necessary. |
| Attributes | The columns that you want to use as the match criteria. The attributes are the columns that are available for selection in a report. From the list of **Available Columns**, you can add the selection criteria for matching data and for aggregating the details (the **Selected Columns**). |
| Sort By/ Group By | The **Sort By** and the **Group By** criteria allow you to organize/segment the data in the report; the sorting and grouping attributes available vary based on the selected data source.<br>The Sort By option specifies the attribute that is used for aggregation. If you do not select an attribute to sort by, the report will return the first N number of results without any aggregation.<br>The Group By option allows you to select an attribute and use it as an anchor for grouping data; all the data in the report is then presented in a set of top 5, 10, 25 or 50 groups. For example, when you select Hour as the group by selection and want the top 25 groups for a 24-hr time period. The results of the report will be generated on an hourly basis over a 24-hr period. The first column in the report will be the hour and then the next set of columns will be the rest of your selected report columns. |

© Palo Alto Networks, Inc.

| Selection | Description |
|---|---|
| | The following example illustrates how the **Selected Columns** and **Sort By**/**Group By** criteria work together when generating reports:<br> 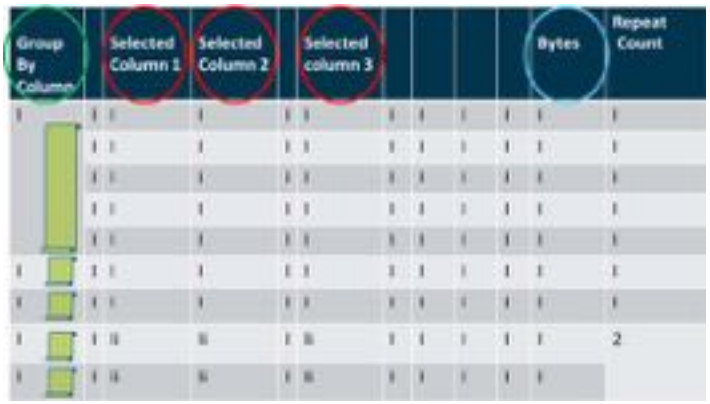 <br>The columns circled in red (above) depict the columns selected, which are the attributes that you match against for generating the report. Each log entry from the data source is parsed and these columns are matched on. If multiple sessions have the same values for the selected columns, the sessions are aggregated and the repeat count (or sessions) is incremented.<br>The column circled in blue indicates the chosen sort order. When the sort order (**Sort By**) is specified, the data is sorted (and aggregated) by the selected attribute.<br>The column circled in green indicates the **Group By** selection, which serves as an anchor for the report. The **Group By** column is used as a match criteria to filter for the top N groups. Then, for each of the top N groups the report enumerates the values for all the other selected columns. |

 © Palo Alto Networks, Inc.

| Selection | Description |
|---|---|

For example, if a report has the following selections

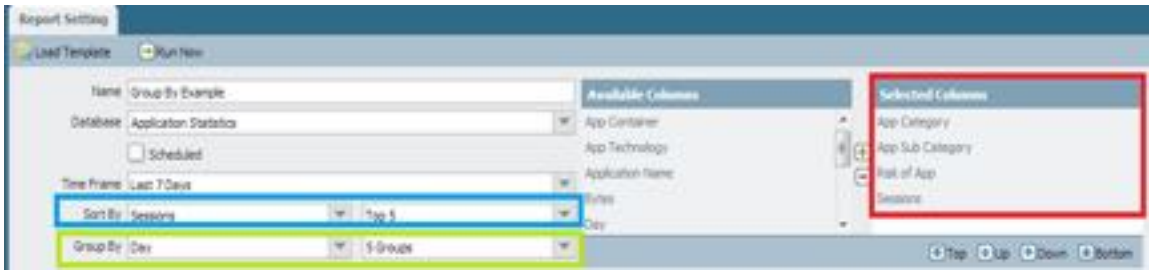


The output will display as follows:

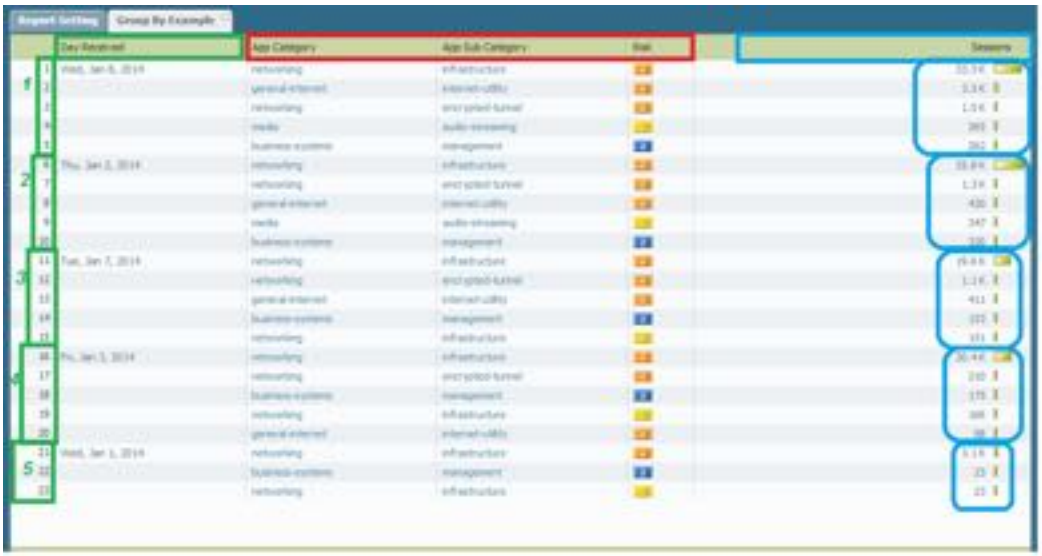

The report is anchored by **Day** and sorted by **Sessions**. It lists the 5 days (**5 Groups**) with maximum traffic in the **Last 7 Days** time frame. The data is enumerated by the **Top 5** sessions for each day for the selected columns—**App Category**, **App Subcategory** and **Risk**.

| | |
|---|---|
| Time Period | The date range for which you want to analyze data. You can define a custom range or select a time period ranging from last 15 minutes to the last 30 days. The reports can be run on demand or scheduled to run at a daily or weekly cadence. |
| Query Builder | The query builder allows you to define specific queries to further refine the selected attributes. It allows you see just what you want in your report using **and** and **or** operators and a match criteria, and then include or exclude data that matches or negates the query in the report. Queries enable you to generate a more focused collation of information in a report. |

## Generate Custom Reports

1. Select **Monitor > Manage Custom Reports**.
2. Click **Add** and then enter a **Name** for the report.

   To base a report on an predefined template, click **Load Template** and choose the template. You can then edit the template and save it as a custom report.
3. Select the database to use for the report.

Each time you create a custom report, a **Log View** report is automatically created. This report show the logs that were used to build the custom report. The log view report uses the same name as the custom report, but appends the phrase (Log View) to the report name.

When creating a report group, you can include the log view report with the custom report. For more information, see Manage Report Groups.

4. Select the **Scheduled** check box to run the report each night. The report is then available for viewing in the **Reports** column on the side.
5. Define the filtering criteria. Select the **Time Frame**, the **Sort By** order, **Group By** preference, and select the columns that must display in the report.
6. (Optional) Select the **Query Builder** attributes, if you want to further refine the selection criteria. To build a report query, specify the following and click **Add**. Repeat as needed to construct the full query.
   - **Connector**—Choose the connector (and/or) to precede the expression you are adding.
   - **Negate**—Select the check box to interpret the query as a negation. If for example, you choose to match entries in the last 24 hours and/or are originating from the untrust zone, the negate option causes a match on entries that are not in the past 24 hours and/or are not from the untrust zone.
   - **Attribute**—Choose a data element. The available options depend on the choice of database.
   - **Operator**—Choose the criterion to determine whether the attribute applies (such as =). The available options depend on the choice of database.
   - **Value**—Specify the attribute value to match.

     For example, the following figure (based on the Traffic Log database) shows a query that matches if the traffic log entry was received in the past 24 hours and is from the "untrust" zone.

7. To test the report settings, select **Run Now**. Modify the settings as required to change the information that is displayed in the report.
8. Click **OK** to save the custom report.


© Palo Alto Networks, Inc.

**Generate Custom Reports (Continued)**

**Examples of Custom Reports**

If we now set up a simple report in which we use the traffic summary database from the last 30 days, and sort the data by the top 10 sessions and these sessions are grouped into 5 groups by day of the week. You would set up the custom report to look like this:

And the PDF output for the report would look as follows:

**Generate Custom Reports (Continued)**

Now, if you want to use the query builder to generate a custom report that represents the top consumers of network resources within a user group, you would set up the report to look like this:

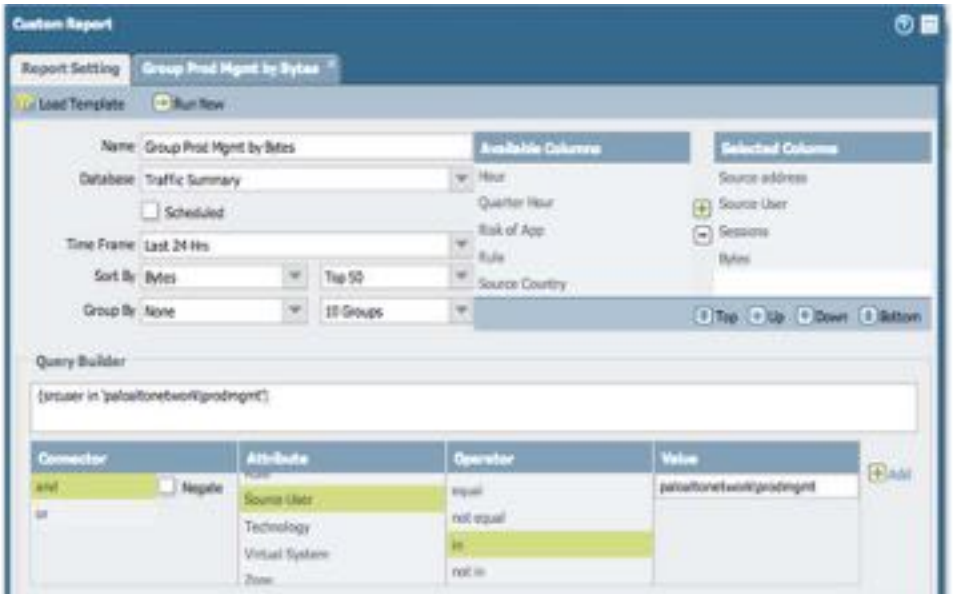


The report would display the top users in the product management user group sorted by bytes, as follows:

 © Palo Alto Networks, Inc.

# Generate Botnet Reports

The botnet report enables you to use heuristic and behavior-based mechanisms to identify potential malware- or botnet-infected hosts in your network. To evaluate botnet activity and infected hosts, the firewall correlates user and network activity data in Threat, URL, and Data Filtering logs with the list of malware URLs in PAN-DB, known dynamic DNS domain providers, and domains registered within the last 30 days. You can configure the report to identify hosts that visited those sites, as well as hosts that communicated with Internet Relay Chat (IRC) servers or that used unknown applications. Malware often use dynamic DNS to avoid IP blacklisting, while IRC servers often use bots for automated functions.

The firewall requires Threat Prevention and URL Filtering licenses to use the botnet report.

- Configure a Botnet Report
- Interpret Botnet Report Output

## Configure a Botnet Report

You can schedule a botnet report or run it on demand. The firewall generates scheduled botnet reports every 24 hours because behavior-based detection requires correlating traffic across multiple logs over that timeframe.

| Configure a Botnet Report | |
|---|---|
| Step 1 Define the types of traffic that indicate possible botnet activity. | 1. Select **Monitor > Botnet** and click **Configuration** on the right side of the page.<br>2. **Enable** and define the **Count** for each type of HTTP Traffic that the report will include.<br>The **Count** values represent the minimum number of events of each traffic type that must occur for the report to list the associated host with a higher confidence score (higher likelihood of botnet infection). If the number of events is less than the **Count**, the report will display a lower confidence score or (for certain traffic types) won't display an entry for the host. For example, if you set the **Count** to three for **Malware URL visit**, then hosts that visit three or more known malware URLs will have higher scores than hosts that visit less than three. For details, see Interpret Botnet Report Output.<br>3. Define the thresholds that determine whether the report will include hosts associated with traffic involving Unknown TCP or Unknown UDP applications.<br>4. Select the **IRC** check box to include traffic involving IRC servers.<br>5. Click **OK** to save the report configuration. |

| Configure a Botnet Report (Continued) | |
|---|---|
| Step 2 Schedule the report or run it on demand. | 1. Click **Report Setting** on the right side of the page.<br>2. Select a time interval for the report in the **Test Run Time Frame** drop-down.<br>3. Select the **No. of Rows** to include in the report.<br>4. (Optional) **Add** queries to the Query Builder to filter the report output by attributes such as source/destination IP addresses, users, or zones.<br>For example, if you know in advance that traffic initiated from the IP address 10.3.3.15 contains no potential botnet activity, you can add `not (addr.src in 10.0.1.35)` as a query to exclude that host from the report output. For details, see Interpret Botnet Report Output.<br>5. Select **Scheduled** to run the report daily or click **Run Now** to run the report immediately.<br>6. Click **OK** and **Commit**. |

## Interpret Botnet Report Output

The botnet report displays a line for each host that is associated with traffic you defined as suspicious when configuring the report. For each host, the report displays a confidence score of 1 to 5 to indicate the likelihood of botnet infection, where 5 indicates the highest likelihood. The scores correspond to threat severity levels: 1 is informational, 2 is low, 3 is medium, 4 is high, and 5 is critical. The firewall bases the scores on:

- **Traffic type**—Certain HTTP traffic types are more likely to involve botnet activity. For example, the report assigns a higher confidence to hosts that visit known malware URLs than to hosts that browse to IP domains instead of URLs, assuming you defined both those activities as suspicious.
- **Number of events**—Hosts that are associated with a higher number of suspicious events will have higher confidence scores based on the thresholds (**Count** values) you define when you Configure a Botnet Report.
- **Executable downloads**—The report assigns a higher confidence to hosts that download executable files. Executable files are a part of many infections and, when combined with the other types of suspicious traffic, can help you prioritize your investigations of compromised hosts.

When reviewing the report output, you might find that the sources the firewall uses to evaluate botnet activity (for example, the list of malware URLs in PAN-DB) have gaps. You might also find that these sources identify traffic that you consider safe. To compensate in both cases, you can add query filters when you Configure a Botnet Report.

 © Palo Alto Networks, Inc.

# Manage PDF Summary Reports

PDF summary reports contain information compiled from existing reports, based on data for the top 5 in each category (instead of top 50). They also contain trend charts that are not available in other reports.

| **Generate PDF Summary Reports** | |
|---|---|
| Step 1 Set up a **PDF Summary Report**. | 1. Select **Monitor > PDF Reports > Manage PDF Summary**.<br>2. Click **Add** and then enter a **Name** for the report.<br>3. Use the drop-down list for each report group and select one or more of the elements to design the PDF Summary Report. You can include a maximum of 18 report elements.<br> 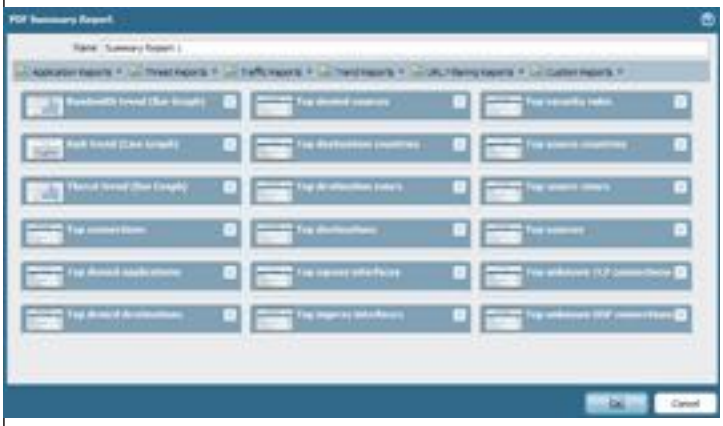 <br>• To remove an element from the report, click the **x** icon or clear the selection from the drop-down for the appropriate report group.<br>• To rearrange the reports, drag and drop the icons to another area of the report.<br>4. Click **OK t**o save the report.<br>5. **Commit** the changes. |

| Generate PDF Summary Reports | |
|---|---|
| Step 2 View the report. | To download and view the PDF Summary Report, see View Reports. 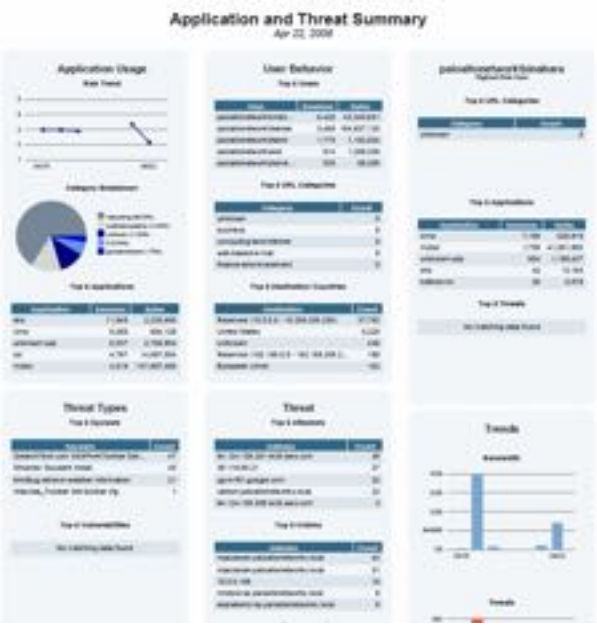  |

 © Palo Alto Networks, Inc.

## Generate User/Group Activity Reports

User/Group activity reports summarize the web activity of individual users or user groups. Both reports include the same information with a couple exceptions—**Browsing Summary by URL Category** and **Browse time calculations** are included in User Activity Reports, but are not included in Group Activity Reports.

User-ID must be configured on the firewall, in order to access the list of user/user groups.

**Generate User/Group Activity Reports**

1. Select **Monitor > PDF Reports > User Activity Report**.
2. Click **Add** and then enter a **Name** for the report.
3. Create the report:
   - For a User Activity Report: Select **User** and enter the **Username** or **IP address** (IPv4 or IPv6) of the user who will be the subject of the report.
   - For Group Activity Report: Select **Group** and select the **Group Name** for which to retrieve user group information in the report.
4. Select the time frame for the report from the drop-down.

   The number of logs that are analyzed in a user activity report is determined by the number of rows defined on the **Max Rows in User Activity Report** on the Logging and Reporting Settings section on **Device >Setup > Management**.
5. Select **Include Detailed Browsing** to include detailed URL logs in the report.

   The detailed browsing information can include a large volume of logs (thousands of logs) for the selected user or user group and can make the report very large.
6. To run the report on demand, click **Run Now**.
7. To save the report, click **OK**. User/Group activity reports cannot be saved on the firewall; to schedule the report for email delivery, see Schedule Reports for Email Delivery.

# Manage Report Groups

Report groups allow you to create sets of reports that the system can compile and send as a single aggregate PDF report with an optional title page and all the constituent reports included.

| Set up Report Groups | |
|---|---|
| Step 1 Set up report groups.<br><br>You must set up a **Report Group** to email report(s). | 1. Create a Server Profile for your email server.<br>2. Define the **Report Group**. A report group can compile predefined reports, PDF Summary reports, custom reports, and Log View report into a single PDF.<br>a. Select **Monitor > Report Group**.<br>b. Click **Add** and then enter a **Name** for the report group.<br>c. (Optional) Select **Title Page** and add a **Title** for the PDF output.<br>d. Select reports from the left column and click **Add** to move each report to the report group on the right.<br><br> 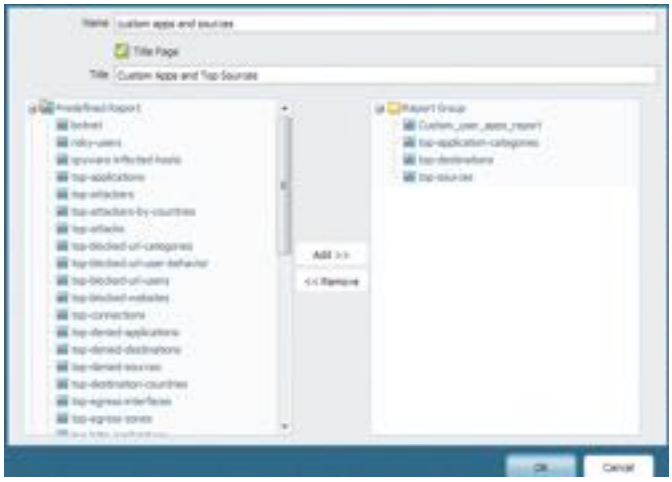 <br><br>The **Log View** report is a report type that is automatically created each time you create a custom report and uses the same name as the custom report. This report will show the logs that were used to build the contents of the custom report.<br><br>To include the log view data, when creating a report group, add your custom report under the **Custom Reports** list and then add the log view report by selecting the matching report name from the **Log View** list. The report will include the custom report data and the log data that was used to create the custom report.<br>e. Click **OK** to save the settings.<br>f. To use the report group, see Schedule Reports for Email Delivery. |


© Palo Alto Networks, Inc.

## Schedule Reports for Email Delivery

Reports can be scheduled for daily delivery or delivered weekly on a specified day. Scheduled reports are executed starting at 2:00 AM, and email delivery starts after all scheduled reports have been generated.

**Schedule Reports for Email Delivery**

1. Select **Monitor > PDF Reports > Email Scheduler**.
2. Select the **Report Group** for email delivery. To set up a report group; see Manage Report Groups.
3. Select the frequency at which to generate and send the report in **Recurrence**.
4. Select the email server profile to use for delivering the reports. To set up an email server profile, see Create a Server Profile for your email server.
5. The **Override Recipient email(s)** allows you to send this report exclusively to the recipients specified in this field. When you add recipients to the **Override Recipient email(s)**, the report is not sent to the recipients configured in the email server profile. Use this option for those occasions when the report is for the attention of someone other than the administrators or recipients defined in the email server profile.

# Syslog Field Descriptions

This is a list of the standard fields for each of the five log types that are forwarded to an external server. For ease of parsing, the comma is the delimiter; each field is a comma-separated value (CSV) string. Fields that are not currently implemented/reserved for future use as tagged as FUTURE_USE.


- Traffic Logs
- Threat Logs
- HIP Match Logs
- Config Logs
- System Logs


## Traffic Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Ingress Interface, Egress Interface, Log Forwarding Profile, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Bytes, Bytes Sent, Bytes Received, Packets, Start Time, Elapsed Time, Category, FUTURE_USE, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Packets Sent, Packets Received.

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Specifies type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of traffic log; values are start, end, drop, and deny<br>• Start—session started<br>• End—session ended<br>• Drop—session dropped before the application is identified and there is no rule that allows the session.<br>• Deny—session dropped after the application is identified and there is a rule to block or no rule that allows the session. |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Source IP (src) | Original session source IP address |
| Destination IP (dst) | Original session destination IP address |
| NAT Source IP (natsrc) | If Source NAT performed, the post-NAT Source IP address |
| NAT Destination IP (natdst) | If Destination NAT performed, the post-NAT Destination IP address |
| Rule Name (rule) | Name of the rule that the session matched |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Source User (srcuser) | Username of the user who initiated the session |
| Destination User (dstuser) | Username of the user to which the session was destined |
| Application (app) | Application associated with the session |
| Virtual System (vsys) | Virtual System associated with the session |
| Source Zone (from) | Zone the session was sourced from |
| Destination Zone (to) | Zone the session was destined to |
| Ingress Interface (inbound_if) | Interface that the session was sourced form |
| Egress Interface (outbound_if) | Interface that the session was destined to |
| Log Forwarding Profile (logset) | Log Forwarding Profile that was applied to the session |
| Session ID (sessionid) | An internal numerical identifier applied to each session |
| Repeat Count (repeatcnt) | Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only |
| Source Port (sport) | Source port utilized by the session |
| Destination Port (dport) | Destination port utilized by the session |
| NAT Source Port (natsport) | Post-NAT source port |
| NAT Destination Port (natdport) | Post-NAT destination port |
| Flags (flags) | 32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value:<br>• 0x80000000—session has a packet capture (PCAP)<br>• 0x02000000—IPv6 session<br>• 0x01000000—SSL session was decrypted (SSL Proxy)<br>• 0x00800000—session was denied via URL filtering<br>• 0x00400000—session has a NAT translation performed (NAT)<br>• 0x00200000—user information for the session was captured via the captive portal (Captive Portal)<br>• 0x00080000—X-Forwarded-For value from a proxy is in the source user field<br>• 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction)<br>• 0x00008000—session is a container page access (Container Page)<br>• 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above<br>• 0x00000800—symmetric return was used to forward traffic for this session |
| Protocol (proto) | IP protocol associated with the session |

| Field Name | Description |
|---|---|
| Action (action) | Action taken for the session; values are allow or deny:<br>• Allow—session was allowed by policy<br>• Deny—session was denied by policy |
| Bytes (bytes) | Number of total bytes (transmit and receive) for the session |
| Bytes Sent (bytes_sent) | Number of bytes in the client-to-server direction of the session<br>Available on all models except the PA-4000 Series |
| Bytes Received (bytes_received) | Number of bytes in the server-to-client direction of the session<br>Available on all models except the PA-4000 Series |
| Packets (packets) | Number of total packets (transmit and receive) for the session |
| Start Time (start) | Time of session start |
| Elapsed Time (elapsed) | Elapsed time of the session |
| Category (category) | URL category associated with the session (if applicable) |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama |
| Source Location (srcloc) | Source country or Internal region for private addresses; maximum length is 32 bytes |
| Destination Location (dstloc) | Destination country or Internal region for private addresses. Maximum length is 32 bytes |
| Packets Sent (pkts_sent) | Number of client-to-server packets for the session<br>Available on all models except the PA-4000 Series |
| Packets Received (pkts_received) | Number of server-to-client packets for the session<br>Available on all models except the PA-4000 Series |

## Threat Logs

**Format**: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Ingress Interface, Egress Interface, Log Forwarding Profile, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Miscellaneous, Threat ID, Category, Severity, Direction, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Content Type, PCAP_id*, Filedigest*, Cloud*

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Specifies type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of threat log; values are URL, virus, spyware, vulnerability, file, scan, flood, data, and WildFire:<br>• url—URL filtering log<br>• virus—virus detection<br>• spyware—spyware detection<br>• vulnerability—vulnerability exploit detection<br>• file—file type log<br>• scan—scan detected via Zone Protection Profile<br>• flood—flood detected via Zone Protection Profile<br>• data—data pattern detected from Data Filtering Profile<br>• wildfire—WildFire log |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Source IP (src) | Original session source IP address |
| Destination IP (dst) | Original session destination IP address |
| NAT Source IP (natsrc) | If source NAT performed, the post-NAT source IP address |
| NAT Destination IP (natdst) | If destination NAT performed, the post-NAT destination IP address |
| Rule Name (rule) | Name of the rule that the session matched |
| Source User (srcuser) | Username of the user who initiated the session |
| Destination User (dstuser) | Username of the user to which the session was destined |
| Application (app) | Application associated with the session |
| Virtual System (vsys) | Virtual System associated with the session |
| Source Zone (from) | Zone the session was sourced from |
| Destination Zone (to) | Zone the session was destined to |
| Ingress Interface (inbound_if) | Interface that the session was sourced from |
| Egress Interface (outbound_if) | Interface that the session was destined to |
| Log Forwarding Profile (logset) | Log Forwarding Profile that was applied to the session |

© Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Session ID (sessionid) | An internal numerical identifier applied to each session |
| Repeat Count (repeatcnt) | Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only |
| Source Port (sport) | Source port utilized by the session |
| Destination Port (dport) | Destination port utilized by the session |
| NAT Source Port (natsport) | Post-NAT source port |
| NAT Destination Port (natdport) | Post-NAT destination port |
| Flags (flags) | 32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value:<br>• 0x80000000—session has a packet capture (PCAP)<br>• 0x02000000—IPv6 session<br>• 0x01000000—SSL session was decrypted (SSL Proxy)<br>• 0x00800000—session was denied via URL filtering<br>• 0x00400000—session has a NAT translation performed (NAT)<br>• 0x00200000—user information for the session was captured via the captive portal (Captive Portal)<br>• 0x00080000—X-Forwarded-For value from a proxy is in the source user field<br>• 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction)<br>• 0x00008000—session is a container page access (Container Page)<br>• 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above<br>• 0x00000800—symmetric return was used to forward traffic for this session |
| Protocol (proto) | IP protocol associated with the session |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Action (action) | Action taken for the session; values are alert, allow, deny, drop, drop-all-packets, reset-client, reset-server, reset-both, block-url.<br>• Alert—threat or URL detected but not blocked<br>• Allow—flood detection alert<br>• Deny—flood detection mechanism activated and deny traffic based on configuration<br>• Drop—threat detected and associated session was dropped<br>• Drop-all-packets—threat detected and session remains, but drops all packets<br>• Reset-client—threat detected and a TCP RST is sent to the client<br>• Reset-server—threat detected and a TCP RST is sent to the server<br>• Reset-both—threat detected and a TCP RST is sent to both the client and the server<br>• Block-url—URL request was blocked because it matched a URL category that was set to be blocked |
| Miscellaneous (misc) | Field with variable length with a maximum of 1023 characters<br>The actual URI when the subtype is URL<br>File name or file type when the subtype is file<br>File name when the subtype is virus<br>File name when the subtype is WildFire |
| Threat ID (threatid) | Palo Alto Networks identifier for the threat. It is a description string followed by a 64-bit numerical identifier in parenthesis for some Subtypes:<br>• 8000 – 8099—scan detection<br>• 8500 – 8599—flood detection<br>• 9999—URL filtering log<br>• 10000 – 19999—sypware phone home detection<br>• 20000 – 29999—spyware download detection<br>• 30000 – 44999—vulnerability exploit detection<br>• 52000 – 52999—filetype detection<br>• 60000 – 69999—data filtering detection<br>• 100000 – 2999999—virus detection<br>• 3000000 – 3999999—WildFire signature feed<br>• 4000000-4999999—DNS Botnet signatures |
| Category (category) | For URL Subtype, it is the URL Category; For WildFire subtype, it is the verdict on the file and is either ‘malicious’ or ‘benign’; For other subtypes the value is ‘any’ |
| Severity (severity) | Severity associated with the threat; Values are informational, low, medium, high, critical |

| Field Name | Description |
|---|---|
| Direction (direction) | Indicates the direction of the attack, *client-to-server* or *server-to-client*<br>• 0—direction of the threat is client to server<br>• 1—direction of the threat is server to client |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama. |
| Source Location (srcloc) | Source country or Internal region for private addresses. Maximum length is 32 bytes. |
| Destination Location (dstloc) | Destination country or Internal region for private addresses. Maximum length is 32 bytes. |
| Content Type (contenttype) | Applicable only when Subtype is URL.<br>Content type of the HTTP response data. Maximum length 32 bytes. |
| New in v6.0!PCAP ID (pcap_id) | Pcap-ID is a 64 bit unsigned integral denoting an ID to correlate threat pcap files with extended pcaps taken as a part of that flow. All threat logs will contain either a pcap_id of 0 (no associated pcap), or an ID referencing the extended pcap file. |
| New in v6.0! File Digest (filedigest) | Only for WildFire subtype; all other types do not use this field<br>The filedigest string shows the binary hash of the file sent to be analyzed by the WildFire service. |
| New in v6.0! Cloud (cloud) | Only for WildFire subtype; all other types do not use this field<br>The cloud string displays the FQDN of either the WildFire appliance (private) or the WildFire cloud (public) from where the file was uploaded for analysis. |

## HIP Match Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source User, Virtual System, Machine name, OS*, Source Address, HIP, Repeat Count, HIP Type, FUTURE_USE, FUTURE_USE, Sequence Number, Action Flags

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of HIP match log; unused |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Source User (srcuser) | Username of the user who initiated the session |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Virtual System (vsys) | Virtual System associated with the HIP match log |
| Machine Name (machinename) | Name of the user's machine |
| New in v6.0! OS | The operating system installed on the user's machine or device (or on the client system) |
| Source Address (src) | IP address of the source user |
| HIP (matchname) | Name of the HIP object or profile |
| Repeat Count (repeatcnt) | Number of times the HIP profile matched |
| HIP Type (matchtype) | Whether the hip field represents a HIP object or a HIP profile |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama |

## Config Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Host, Virtual System, Command, Admin, Client, Result, Configuration Path, Sequence Number, Action Flags

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of configuration log; unused |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Host (host) | Host name or IP address of the client machine |
| Virtual System (vsys) | Virtual System associated with the configuration log |
| Command (cmd) | Command performed by the Admin; values are add, clone, commit, delete, edit, move, rename, set, validate. |
| Admin (admin) | Username of the Administrator performing the configuration |
| Client (client) | Client used by the Administrator; values are Web and CLI |
| Result (result) | Result of the configuration action; values are Submitted, Succeeded, Failed, and Unauthorized |

| Field Name | Description |
|---|---|
| Configuration Path (path) | The path of the configuration command issued; up to 512 bytes in length |
| Sequance Number (seqno) | A 64bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama. |

## System Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Virtual System, Event ID, Object, FUTURE_USE, FUTURE_USE, Module, Severity, Description, Sequence Number, Action Flags

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Virtual System (vsys) | Virtual System associated with the configuration log |
| Event ID (eventid) | String showing the name of the event |
| Object (object) | Name of the object associated with the system event |
| Module (module) | This field is valid only when the value of the Subtype field is general. It provides additional information about the sub-system generating the log; values are general, management, auth, ha, upgrade, chassis |
| Severity (severity) | Severiry associated with the event; values are informational, low, medium, high, critical |
| Description (opaque) | Detailed description of the event, up to a maximum of 512 bytes |
| Sequence Number (seqno) | A 64bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama |

 © Palo Alto Networks, Inc.

## Syslog Severity

The syslog severity is set based on the log type and contents.

| Log Type/Severity | Syslog Severity |
|---|---|
| Traffic | Info |
| Config | Info |
| Threat/System—Informational | Info |
| Threat/System—Low | Notice |
| Threat/System—Medium | Warning |
| Threat/System—High | Error |
| Threat/System—Critical | Critical |

## Custom Log/Event Format

To facilitate the integration with external log parsing systems, the firewall allows you to customize the log format; it also allows you to add custom *Key: Value* attribute pairs. Custom message formats can be configured under **Device > Server Profiles > Syslog > Syslog Server Profile > Custom Log Format**.

To achieve ArcSight Common Event Format (CEF) compliant log formatting, refer to the CEF Configuration Guide.

## Escape Sequences

Any field that contains a comma or a double-quote is enclosed in double quotes. Furthermore, if a double-quote appears inside a field it is escaped by preceding it with another double-quote. To maintain backward compatibility, the Misc field in threat log is always enclosed in double-quotes.

 © Palo Alto Networks, Inc.

# User-ID

User Identification (User-ID) is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses. The following sections describe the Palo Alto Networks User-ID feature and provide instructions on setting up user- and group-based access:


- User-ID Overview
- User-ID Concepts
- Enable User-ID
- Map Users to Groups
- Map IP Addresses to Users
- Enable User- and Group-Based Policy
- Verify the User-ID Configuration


© Palo Alto Networks, Inc.

# User-ID Overview

User-ID seamlessly integrates Palo Alto Networks firewalls with a range of enterprise directory and terminal services offerings, enabling you to tie application activity and security policies to users and groups—not just IP addresses. In addition, with User-ID enabled, the Application Command Center (ACC), App-Scope, reports, and logs all include usernames in addition to user IP addresses.

The Palo Alto Networks next-generation firewall supports monitoring of the following enterprise services:

- Microsoft Active Directory
- LDAP
- Novell eDirectory
- Citrix Metaframe Presentation Server or XenApp
- Microsoft Terminal Services

To be able to create policy based on user and group, the firewall must have a list of all available users and their corresponding group mappings that you can select from when defining your policies. It gets this Group Mapping information by connecting directly to your LDAP directory server.

To be able to enforce the user- and group-based policies, the firewall must be able to map the IP addresses in the packets it receives to usernames. User-ID provides many mechanisms to get these IP address to username mappings. For example, it uses agents to monitor server logs for login events and/or probe clients, and/or listen for syslog messages from authenticating services. To identify mappings for IP addresses were not mapped using one of the agent mechanisms, you can configure the firewall to redirect HTTP requests to a captive portal login. You can tailor the mechanisms you use for User Mapping to suit your environment, and you can even use different mechanisms at different sites.

 © Palo Alto Networks, Inc.

**Figure: User-ID**

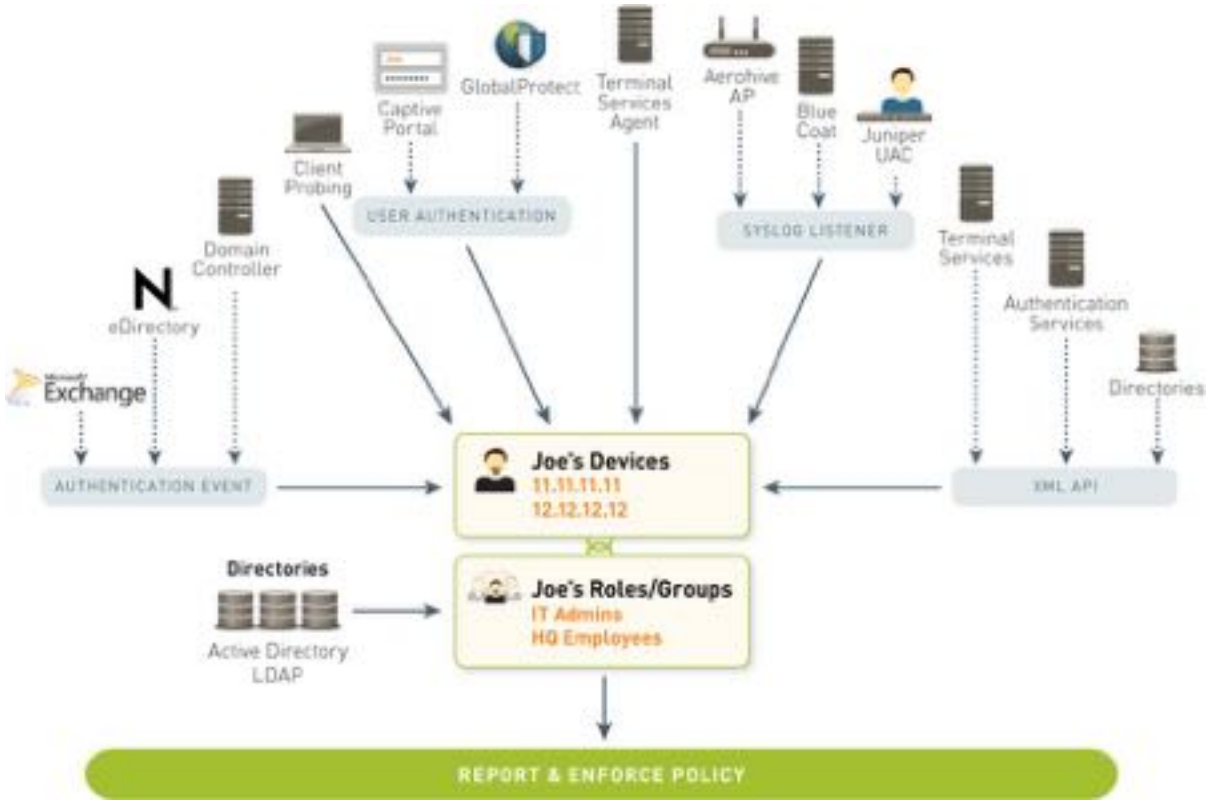


Proceed to User-ID Concepts for information on how User-ID works and Enable User-ID for instructions on setting up User-ID on the firewall.

# User-ID Concepts

▲ Group Mapping
▲ User Mapping

## Group Mapping

In order to define security policies based on user or group, the firewall must retrieve the list of groups and the corresponding list of members from your directory server. To enable this functionality, you must create an LDAP server profile that instructs the firewall how to connect and authenticate to the server and how to search the directory for the user and group information. After you connect to the LDAP server and configure the group mapping functionality for user identification, you will be able to select users or groups when defining your security policies. The firewall supports a variety of LDAP directory servers, including Microsoft Active Directory (AD), Novell eDirectory, and Sun ONE Directory Server.

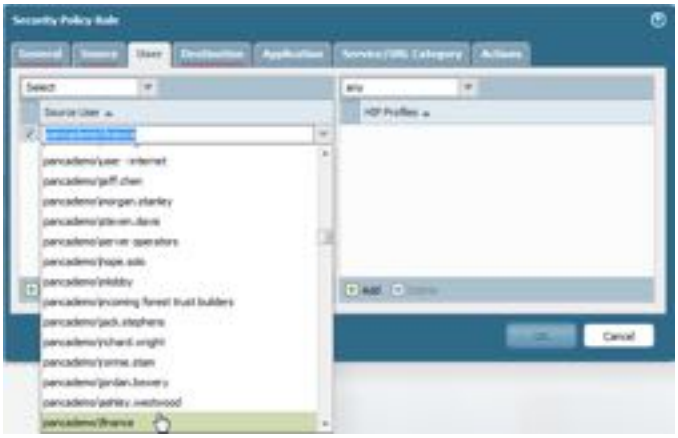

You can then define policies based on group membership rather than on individual users for simplified administration. For example, the following security policy allows access to specific internal applications based on group membership:

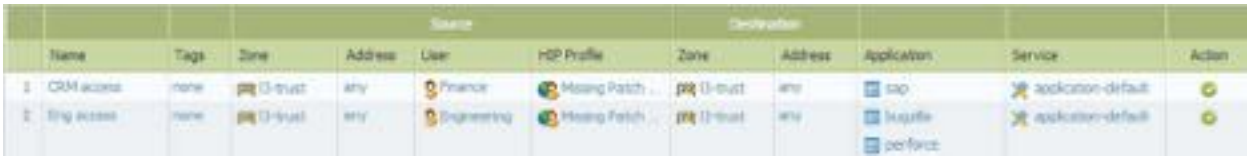

## User Mapping

Having the names of the users and groups is only one piece of the puzzle. The firewall also needs to know which IP addresses map to which users so that security policies can be enforced appropriately. Figure: User-ID illustrates the different methods that are used to identify users and groups on your network and shows how user mapping and group mapping work together to enable user- and group-based security enforcement and visibility.

The following topics describe the different methods of user mapping:

 © Palo Alto Networks, Inc.

- Server Monitoring
- Client Probing
- Port Mapping
- Syslog
- Captive Portal
- GlobalProtect
- User-ID XML API

## Server Monitoring

With server monitoring a User-ID agent—either a Windows-based agent running on a domain server in your network, or the integrated PAN-OS User-ID agent running on the firewall—monitors the security event logs for specified Microsoft Exchange Servers, domain controllers, or Novell eDirectory servers for login events. For example, in an AD environment, you can configure the User-ID agent to monitor the security logs for Kerberos ticket grants or renewals, Exchange server access (if configured), and file and print service connections. Keep in mind that in order for these events to be recorded in the security log, the AD domain must be configured to log successful account login events. In addition, because users can log in to any of the servers in the domain, you must set up server monitoring for all servers in order to capture all user login events.

Because server monitoring requires very little overhead and because the majority of users can generally be mapped using this method, it is recommended as the base user mapping method for most User-ID deployments. See Configure User Mapping Using the Windows User-ID Agent or Configure User Mapping Using the PAN-OS Integrated User-ID Agent for details.

## Client Probing

In a Microsoft Windows environment, you can configure the User-ID agent to probe client systems using Windows Management Instrumentation (WMI). The Windows-based User-ID agent can also perform NetBIOS probing (not supported on the PAN-OS integrated User-ID agent). Probing is particularly useful in environments with a high IP address turnover because changes will be reflected on the firewall more quickly, enabling more accurate enforcement of user-based policies. However, if the correlation between IP addresses and users is fairly static, you probably do not need to enable client probing. Because probing can generate a large amount of network traffic (based on the total number of mapped IP addresses), the agent that will be initiating the probes should be located as close as possible to the end clients.

If probing is enabled, the agent will probe each learned IP address periodically (every 20 minutes by default, but this is configurable) to verify that the same user is still logged in. In addition, when the firewall encounters an IP address for which it has no user mapping it will send the address to the agent for an immediate probe.

See Configure User Mapping Using the Windows User-ID Agent or Configure User Mapping Using the PAN-OS Integrated User-ID Agent for details.

© Palo Alto Networks, Inc.

## Port Mapping

In environments with multi-user systems—such as Microsoft Terminal Server or Citrix environments—many users share the same IP address. In this case, the user-to-IP address mapping process requires knowledge of the source port of each client. To perform this type of mapping, you must install the Palo Alto Networks Terminal Services Agent on the Windows/Citrix terminal server itself to intermediate the assignment of source ports to the various user processes. For terminal servers that do not support the Terminal Services Agent, such as Linux terminal servers, you can use the XML API to send user mapping information from login and logout events to User-ID. See Configure User Mapping for Terminal Server Users for configuration details.

## Syslog

In environments with existing network services that authenticate users—such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms—the firewall User-ID agent (either the Windows agent or the PAN-OS integrated agent on the firewall) can listen for authentication syslog messages from those services. Syslog filters, which are provided by a content update (integrated User-ID agent only) or configured manually, allow the User-ID agent to parse and extract usernames and IP addresses from authentication syslog events generated by the external service, and add the information to the User-ID IP address to username mappings maintained by the firewall. See Configure User-ID to Receive User Mappings from a Syslog Sender for configuration details.

**Figure: User-ID Integration with Syslog**

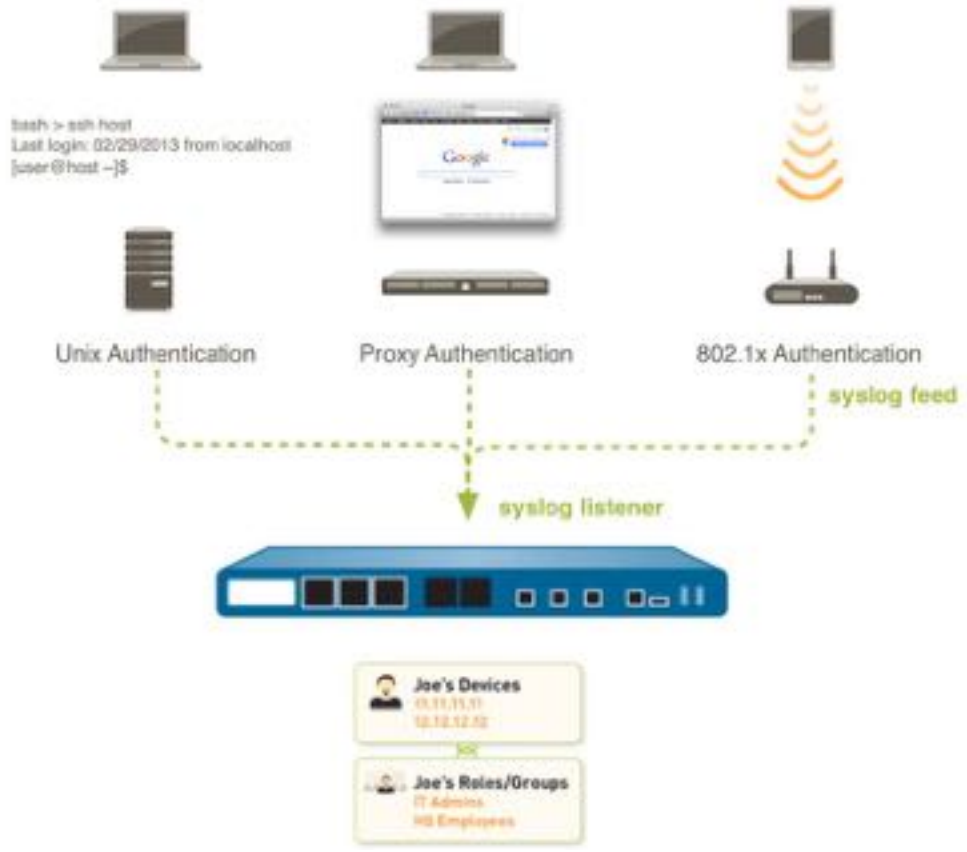


 © Palo Alto Networks, Inc.

## Captive Portal

If the firewall or the User-ID agent are unable to map an IP address to a user—for example if the user is not logged in or is using an operating system such as Linux that is not supported by your domain servers—you can configure Captive Portal. When configured, any web traffic (HTTP or HTTS) matching your Captive Portal policy requires user authentication, either transparently via an NT LAN Manager (NTLM) challenge to the browser, or actively by redirecting the user to a web authentication form for authentication against a RADIUS, LDAP, Kerberos, or local authentication database or using client certificate authentication. See Map IP Addresses to User Names Using Captive Portal for details.

## GlobalProtect

For mobile or roaming users, the GlobalProtect client provides the user mapping information to the firewall directly. In this case, every GlobalProtect user has an agent or app running on the client that requires the user to enter login credentials for VPN access to the firewall. This login information is then added to the User-ID user mapping table on the firewall for visibility and user-based security policy enforcement. Because GlobalProtect users must authenticate to gain access to the network, the IP address to username mapping is explicitly known. This is the best solution in sensitive environments where you must be certain of who a user is in order to allow access to an application or service. For more information on setting up GlobalProtect, refer to the GlobalProtect Administrator's Guide.

## User-ID XML API

For other types of user access that cannot be mapped using any of the standard user mapping methods or Captive Portal—for example, to add mappings of users connecting from a third-party VPN solution or users connecting to a 802.1x enabled wireless network—you can use the User-ID XML API to capture login events and send them to the User-ID agent or directly to the firewall. See Send User Mappings to User-ID Using the XML API for details.

© Palo Alto Networks, Inc.

# Enable User-ID

You must complete the following tasks to set up the firewall to user users and groups in policy enforcement, logging, and reporting:

- Map Users to Groups
- Map IP Addresses to Users
- Enable User- and Group-Based Policy
- Verify the User-ID Configuration

 © Palo Alto Networks, Inc.

# Map Users to Groups

Use the following procedure to connect to your LDAP directory to enable the firewall to retrieve user-to-group mapping information:

Best practices for group mapping in an Active Directory environment:

- If you have a single domain, you only need one LDAP server profile that connects the firewall to the domain controller with the best connectivity. You can add additional domain controllers for fault tolerance.
- If you have multiple domains and/or multiple forests, you must create a server profile to connect to a domain server in each domain/forest. Take steps to ensure unique usernames in separate forests.
- If you have Universal Groups, create a server profile to connect to the Global Catalogue server.

## Map Users to Groups

Step 1 Create an LDAP Server Profile that specifies how to connect to the directory servers you want the firewall to use to obtain group mapping information.

1. Select **Device > Server Profiles > LDAP**.
2. Click **Add** and then enter a **Name** for the profile.
3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.
4. Click **Add** to add a new LDAP server entry and then enter a **Server** name to identify the server (1-31 characters) and the IP **Address** and **Port** number the firewall should use to connect to the LDAP server (default=389 for LDAP; 636 for LDAP over SSL). You can add up to four LDAP servers to the profile, however, all the servers you add to a profile must be of the same type. For redundancy you should add at least two servers.

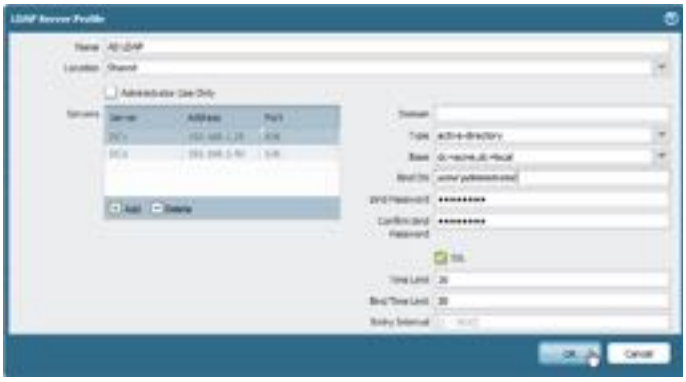

5. Enter the LDAP **Domain** name to prepend to all objects learned from the server. The value you enter here depends on your deployment:
   - If you are using Active Directory, you must enter the NetBIOS domain name; NOT a FQDN (for example, enter `acme`, not `acme.com`). Note that if you need to collect data from multiple domains you will need to create separate server profiles.
   - If you are using a global catalog server, leave this field blank.
6. Select the **Type** of LDAP server you are connecting to. The correct LDAP attributes in the group mapping settings will automatically be populated based on your selection. However, if you have customized your LDAP schema you may need to modify the default settings.
7. In the **Base** field, select the DN that corresponds to the point in the LDAP tree where you want the firewall to begin its search for user and group information.
8. Enter the authentication credentials for binding to the LDAP tree in the **Bind DN**, **Bind Password**, and **Confirm Bind Password** fields. The Bind DN can be in either User Principal Name (UPN) format (for example, `administrator@acme.local`) or it can be a fully qualified LDAP name (for example, `cn=administrator,cn=users,dc=acme,dc=local`).
9. If you want the firewall to communicate with the LDAP server(s) over a secure connection, select the SSL check box. If you enable SSL, make sure that you have also specified the appropriate port number.
10. Click **OK** to save the profile.

**Map Users to Groups (Continued)**

Step 2 Add the LDAP server profile to the User-ID Group Mapping configuration.

1. Select **Device > User Identification > Group Mapping Settings** and click **Add**.
2. Select the **Server Profile** you created in Step 1.
3. Make sure the **Enabled** check box is selected.
4. (Optional) If you want to limit which groups are displayed within security policy, select the **Group Include List** tab and then browse through the LDAP tree to locate the groups you want to be able to use in policy. For each group you want to include, select it in the **Available Groups** list and click the add icon to move it to the **Included Groups** list. Repeat this step for every group you want to be able to use in your policies.
5. Click **OK** to save the settings.

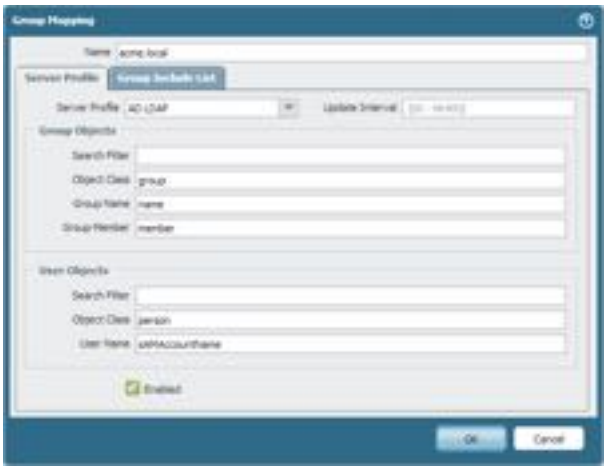

Step 3 **Commit** the configuration.


© Palo Alto Networks, Inc.

# Map IP Addresses to Users

The tasks you need to perform to map IP addresses to usernames depends on the type and location of the client systems on your network. Complete as many of the following tasks as necessary to enable mapping of your client systems:

- To map users as they log in to your Exchange servers, domain controllers, or eDirectory servers, or Windows clients that can be directly probed you must configure a User-ID agent to monitor the server logs and/or probe client systems. You can either install the standalone Windows User-ID agent on one or more member servers in the domain that contains the servers and clients to be monitored (see Configure User Mapping Using the Windows User-ID Agent) or you can configure the on-firewall User-ID agent that is integrated with PAN-OS (Configure User Mapping Using the PAN-OS Integrated User-ID Agent). For guidance as to which agent configuration is appropriate for your network and the number and placements of agents that are required, refer to Architecting User Identification Deployments.
- If you have clients running multi-user systems such as Microsoft Terminal Server or Citrix Metaframe Presentation Server or XenApp, see Configure the Palo Alto Networks Terminal Server Agent for User Mapping for instructions on how to install and configure the agent on a Windows server. If you have a multi-user system that is not running on Windows, you can use the User-ID XML API to send IP address to username mappings directly to the firewall. See Retrieve User Mappings from a Terminal Server Using the User-ID XML API.
- To obtain user mappings from existing network services that authenticate users, such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms, configure the User-ID agent (either the Windows agent or the agentless user mapping feature on the firewall) to listen for authentication syslog messages from those services. See Configure User-ID to Receive User Mappings from a Syslog Sender.
- If you have users with client systems that are not logged into your domain servers—for example, users running Linux clients that do not log in to the domain—see Map IP Addresses to User Names Using Captive Portal.
- For other clients that you are unable to map using the previous methods, you can use the User-ID XML API to add user mappings directly to the firewall. See Send User Mappings to User-ID Using the XML API.
- Because policy is local to each firewall, each firewall must have a current list of IP address to username mappings in order to accurately enforce security policy by group or user. However, you can configure one firewall to collect all the user mappings and distribute them to the other firewalls. For details, see Configure a Firewall to Share User Mapping Data with Other Firewalls.

© Palo Alto Networks, Inc.

# Configure User Mapping Using the Windows User-ID Agent

In most cases, the majority of your network users will have logins to your monitored domain services. For these users, the Palo Alto Networks User-ID agent monitors the servers for login events and performs the IP address to user mapping. The way you configure the User-ID agent depends on the size of your environment and the location of your domain servers. As a best practice, you should locate your User-ID agents near your monitored servers (that is, the monitored servers and the Windows User-ID agent should not be across a WAN link from each other). This is because most of the traffic for user mapping occurs between the agent and the monitored server, with only a small amount of traffic—the delta of IP address mappings since the last update—from the agent to the firewall.

The following topics describe how to install and configure the User-ID Agent and how to configure the firewall to retrieve user mapping information from the agent:

- Install the User-ID Agent
- Configure the User-ID Agent for User Mapping

## Install the User-ID Agent

The following procedure shows how to install the User-ID agent on a member server in the domain and set up the service account with the required permissions. If you are upgrading, the installer will automatically remove the older version, however, it is a good idea to back up the config.xml file before running the installer.

For information about the system requirements for installing the Windows-based User-ID agent and for information on the supported server OS versions are supported, refer to "Operating System (OS) Compatibility User-ID Agent" in the User-ID Agent Release Notes, which are available on the Palo Alto Networks Software Updates page.

| Install the Windows User-ID Agent | |
|---|---|
| Step 1 Decide where to install the User-ID agent(s).<br><br>The User-ID agent queries the Domain Controller and Exchange server logs using Microsoft Remote Procedure Calls (MSRPCs), which require a complete transfer of the entire log at each query. Therefore, you should always install one or more User-ID agents at each site that has servers to be monitored.<br><br> For more detailed information on where to install User-ID agents, refer to Architecting User Identification (User-ID) Deployments. | • You must install the User-ID agent on a system running one of the following OS versions (32-bit and 64-bit are both supported):<br>  • Microsoft Windows XP/Vista/7<br>  • Microsoft Windows Server 2003/2008<br>• Make sure the system you plan to install the User-ID agent on is a member of the domain that the servers it will be monitoring belong to.<br>• As a best practice, install the User-ID agent close to the servers it will be monitoring (there is more traffic between the User-ID agent and the monitored servers than there is between the User-ID agent and the firewall, so locating the agent close to the monitored servers optimizes bandwidth usage).<br>• To ensure the most comprehensive mapping of users, you must monitor all servers that contain user login information. You may need to install multiple User-ID agents to efficiently monitor all of your resources. |

 © Palo Alto Networks, Inc.

**Install the Windows User-ID Agent (Continued)**

| | |
|---|---|
| Step 2 Download the User-ID agent installer.<br><br>As a best practice, install the User-ID agent version that is the same as the PAN-OS version running on the firewalls. | 1. Log in to Palo Alto Networks Support site.<br>2. Select **Software Updates** from the Manage Devices section.<br>3. Scroll to the User Identification Agent section of the screen and **Download** the version of the User-ID agent you want to install.<br>4. Save the `UaInstall-x.x.x-xx.msi` file on the system(s) where you plan to install the agent. |
| Step 3 Run the installer as an administrator.<br> | 1. To launch a command prompt as an administrator, click Start and right-click **Command Prompt** and then select **Run as administrator**.<br>2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop you would enter the following:<br>`C:\Users\administrator.acme>cd Desktop`<br>`C:\Users\administrator.acme\Desktop>UaInstall-6.0.0-1.msi`<br>3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to the `C:\Program Files (x86)\Palo Alto Networks\User-ID Agent` folder, but you can **Browse** to a different location.<br>4. When the installation completes, **Close** the setup window. |
| Step 4 Launch the User-ID Agent application.<br> 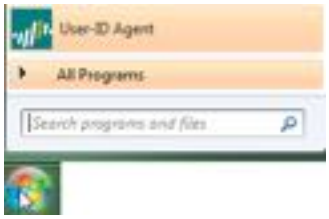 | 1. Click Start and select **User-ID Agent**. |
| Step 5 (Optional) Change the service account that the User-ID agent uses to log in.<br> | By default, the agent uses the administrator account used to install the .msi file. However, you may want to switch this to a restricted account as follows:<br>1. Select **User Identification > Setup** and click **Edit**.<br>2. Select the **Authentication** tab and enter the service account name that you want the User-ID agent to use in the **User name for Active Directory** field.<br>3. Enter the **Password** for the specified account. |

| Install the Windows User-ID Agent (Continued) | |
|---|---|
| Step 6 (Optional) Assign account permissions to the installation folder.<br><br>You only need to perform this step if the service account you configured for the User-ID agent is not a member of the administrators group for the domain or a member of both the Server Operators and the Event Log Readers groups. | 1. Give the service account permissions to the installation folder:<br>a. From the Windows Explorer, navigate to `C:\Program Files\Palo Alto Networks` and right-click the folder and select **Properties**.<br>b. On the **Security** tab, **Add** the User-ID agent service account and assign it permissions to **Modify**, **Read & execute**, **List folder contents**, and **Read** and then click **OK** to save the account settings.<br>2. Give the service account permissions to the User-ID Agent registry sub-tree:<br>a. Run `regedit32` and navigate to the Palo Alto Networks sub-tree in one of the following locations:<br>– **32-bit systems**—`HKEY_LOCAL_MACHINE\Software\Palo Alto Networks`<br>– **64-bit systems**—`HKEY_LOCAL_MACHINE\Software\WOW6432Node\Palo Alto Networks`<br>b. Right-click the Palo Alto Networks node and select **Permissions**.<br>c. Assign the User-ID service account **Full Control** and then click **OK** to save the setting.<br>3. On the domain controller, add the service account to the builtin groups to enable privileges to read the security log events (Event Log Reader group) and open sessions (Server Operator group):<br>a. Run the MMC and Launch the Active Directory Users and Computers snap-in.<br>b. Navigate to the Builtin folder for the domain and then right-click each group you need to edit (Event Log Reader and Server Operator) and select **Add to Group** to open the properties dialog.<br>c. Click **Add** and enter the name of the service account that you configured the User-ID service to use and then click **Check Names** to validate that you have the proper object name.<br>d. Click **OK** twice to save the settings. |

## Configure the User-ID Agent for User Mapping

The Palo Alto Networks User-ID agent is a Windows service that connects to servers on your network—for example, Active Directory servers, Microsoft Exchange servers, and Novell eDirectory servers—and monitors the logs for login events. The agent uses this information to map IP addresses to usernames. Palo Alto Networks firewalls connect to the User-ID agent to retrieve this user mapping information, enabling visibility into user activity by username rather than IP address and enables user- and group-based security enforcement.

 © Palo Alto Networks, Inc.

For information about the server OS versions supported by the User-ID agent, refer to "Operating System (OS) Compatibility User-ID Agent" in the *User-ID Agent Release Notes*, which are available on the Palo Alto Networks Software Updates page.

| Map IP Addresses to Users Using the User-ID Agent | | |
|---|---|---|
| Step 1 | Launch the User-ID Agent application. | 1. Select **User-ID Agent** from the Windows Start menu.<br> |
| Step 2 | Define the servers the User-ID agent should monitor to collect IP address to user mapping information.<br>The User-ID agent can monitor up to 100 servers and listen for syslog messages from up to 100 syslog senders.<br>Keep in mind that in order to collect all of the required mappings, you must connect to all servers that your users log in to in order to monitor the security log files on all servers that contain login events. | 1. Select **User Identification > Discovery**.<br>2. In the **Servers** section of the screen, click **Add**.<br>3. Enter a **Name** and **Server Address** for the server to be monitored. The network address can be a FQDN or an IP address.<br>4. Select the **Server Type** (**Microsoft Active Directory**, **Microsoft Exchange**, **Novell eDirectory**, or **Syslog Sender**) and then click **OK** to save the server entry. Repeat this step for each server to be monitored.<br>5. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Auto Discover**.<br>The auto-discovery locates domain controllers in the local domain only; you must manually add Exchange servers, eDirectory servers, and syslog senders.<br>6. (Optional) To tune the frequency at which the firewall polls configured servers for mapping information, select **User Identification > Setup** and **Edit** the Setup section. On the **Server Monitor** tab, modify the value in the **Server Log Monitor Frequency (seconds)** field. As a best practice, you should increase the value in this field to 5 seconds in environments with older Domain Controllers or high-latency links. Click **OK** to save the changes. |

| Map IP Addresses to Users Using the User-ID Agent (Continued) | | |
|---|---|---|
| Step 3 | (Optional) If you configured the agent to connect to a Novell eDirectory server, you must specify how the agent should search the directory. | 1. Select **User Identification > Setup** and click **Edit** in the Setup section of the window.<br>2. Select the **eDirectory** tab and then complete the following fields:<br>• **Search Base**—The starting point or root context for agent queries, for example: `dc=domain1, dc=example, dc=com`.<br>• **Bind Distinguished Name**—The account to use to bind to the directory, for example: `cn=admin, ou=IT, dc=domain1, dc=example, dc=com`.<br>• **Bind Password**—The bind account password. The agent saves the encrypted password in the configuration file.<br>• **Search Filter**—The search query for user entries (default is `objectClass=Person`).<br>• **Server Domain Prefix**—A prefix to uniquely identify the user. This is only required if there are overlapping name spaces, such as different users with the same name from two different directories.<br>• **Use SSL**—Select the check box to use SSL for eDirectory binding.<br>• **Verify Server Certificate**—Select the check box to verify the eDirectory server certificate when using SSL. |
| Step 4 | (Optional) Enable client probing.<br>Client probing is useful in environments where IP addresses are not tightly bound to users because it ensures that previously mapped addresses are still valid. However, as the total number of learned IP addresses grows, so does the amount of traffic generated. As a best practice, only enable probing on network segments where IP address turnover is high.<br>For more details on the placement of User-ID agents using client probing, refer to Architecting User Identification (User-ID) Deployments. | 1. On the **Client Probing** tab, select the **Enable WMI Probing** check box and/or the **Enable NetBIOS Probing** check box.<br>2. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client.<br> For NetBIOS probing to work effectively, each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled. WMI probing is always preferred over NetBIOS whenever possible. |
| Step 5 | Save the configuration. | Click **OK** to save the User-ID agent setup settings and then click **Commit** to restart the User-ID agent and load the new settings. |

 © Palo Alto Networks, Inc.

| Map IP Addresses to Users Using the User-ID Agent (Continued) | | |
|---|---|---|
| Step 6 | (Optional) Define the set of users for which you do not need to provide IP address to user name mappings, such as kiosk accounts.<br><br>You can also use the `ignore-user` list to identify users whom you want to force to authenticate using Captive Portal. | Create an `ignore_user_list.txt` file and save it to the User-ID Agent folder on the domain server where the agent is installed.<br><br>List of the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Each user account name must be on a separate line. For example:<br>`SPAdmin`<br>`SPInstall`<br>`TFSReport` |
| Step 7 | Configure the firewalls to connect to the User-ID agent. | Complete the following steps on each firewall you want to connect to the User-ID agent to receive user mappings:<br>1. Select **Device > User Identification > User-ID Agents** and click **Add**.<br>2. Enter a **Name** for the User-ID agent.<br>3. Enter the IP address of the Windows **Host** on which the User-ID Agent is installed.<br>4. Enter the **Port** number on which the agent will be listening for user mapping requests. This value must match the value configured on the User-ID agent. By default, the port is set to 5007 on the firewall and on newer versions of the User-ID agent. However, some older User-ID agent versions use port 2010 as the default.<br>5. Make sure that the configuration is **Enabled** and then click **OK**.<br>6. **Commit** the changes.<br>7. Verify that the **Connected status** displays as connected. |
| Step 8 | Verify that the User-ID agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent. | 1. Launch the User-ID agent and select **User Identification**.<br>2. Verify that the agent status shows **Agent is running**. If the Agent is not running, click **Start**.<br>3. To verify that the User-ID agent can connect to monitored servers, make sure the Status for each Server is **Connected**.<br>4. To verify that the firewalls can connect to the User-ID agent, make sure the Status for each of the Connected Devices is **Connected**.<br>5. To verify that the User-ID agent is mapping IP addresses to usernames, select **Monitoring** and make sure that the mapping table is populated. You can also **Search** for specific users, or **Delete** user mappings from the list. |

## Configure User Mapping Using the PAN-OS Integrated User-ID Agent

The following procedure shows how to configure the PAN-OS integrated agent on the firewall for user mapping. The integrated User-ID agent performs the same tasks as the Windows-based agent with the exception of NetBIOS client probing (WMI probing is supported).

| Map IP Addresses to Users Using the Integrated User-ID Agent | | |
|---|---|---|
| Step 1 | Create an Active Directory (AD) account for the firewall agent that has the privilege levels required to log in to each service or host you plan to monitor to collect user mapping data. | • **Windows 2008 or later domain servers**—Add the account to the Event Log Readers group. If you are using the on-device User-ID agent, the account must also be a member of the Distributed COM Users Group.<br>• **Windows 2003 domain servers**—Assign Manage Auditing and Security Logs permissions through group policy.<br>• **WMI probing**—Make sure the account has rights to read the CIMV2 namespace; by default, Domain Administrator and Server Operator accounts have this permission.<br>• **NTLM authentication**—Because the firewall must join the domain if you are using Captive Portal NTLM authentication with an on-device User-ID agent, the Windows account you create for NTLM access must have administrative privileges. Note that due to AD restrictions on virtual systems running on the same host, if you have configured multiple virtual systems, only vsys1 will be able to join the domain. |

 © Palo Alto Networks, Inc.

**Map IP Addresses to Users Using the Integrated User-ID Agent (Continued)**

| | |
|---|---|
| Step 2 Define the servers the firewall should monitor to collect IP address to user mapping information. You can define entries for up to 100 Microsoft Active Directory, Microsoft Exchange, or Novell eDirectory servers on your network.<br><br>Keep in mind that in order to collect all of the required mappings, you must connect to all servers that your users log in to so that the firewall can monitor the security log files on all servers that contain login events. | 1. Select **Device > User Identification > User Mapping**.<br>2. In the Server Monitoring section of the screen, click **Add**.<br>3. Enter a **Name** and **Network Address** for the server. The network address can be a FQDN or an IP address.<br>4. Select the **Type** of server.<br>5. Make sure the **Enabled** check box is selected and then click **OK**<br>6. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Discover**.<br> <br>The auto-discovery feature is for domain controllers only; you must manually add any Exchange servers or eDirectory servers you want to monitor.<br>7. (Optional) Specify the frequency at which the firewall polls Windows servers for mapping information. This is the interval between the end of the last query and the start of the next query.<br>If the query load is high, the observed delay between queries might significantly exceed the specified frequency.<br>a. In the **Palo Alto Networks User ID Agent Setup** section, click the Edit icon.<br>b. Select the **Server Monitor** tab and specify the **Server Log Monitor Frequency** in seconds (default is 2, range is 1-3600).<br>As a best practice, increase the value in this field to 5 seconds in environments with older domain controllers or high-latency links.<br>c. Click **OK** to save the changes. |
| Step 3 Set the domain credentials for the account the firewall will use to access Windows resources. This is required for monitoring Exchange servers and domain controllers as well as for WMI probing. | 1. Click the Edit icon in the **Palo Alto Networks User ID Agent Setup** section of the screen.<br>2. On the **WMI Authentication** tab, enter the **User Name** and **Password** for the account that will be used to probe the clients and monitor servers. Enter the user name using the domain\username syntax.<br> |

| Map IP Addresses to Users Using the Integrated User-ID Agent (Continued) | | |
|---|---|---|
| Step 4 | (Optional) Enable WMI probing.<br> The on-device agent does not support NetBIOS probing; it is supported on the Windows-based User-ID agent only. | 1. On the **Client Probing** tab, select the **Enable Probing** check box.<br>2. (Optional) If necessary, modify the **Probe Interval** (in minutes) to ensure it is long enough for the User-ID agent to probe all the learned IP addresses (default is 20, range is 1-1440). This is the interval between the end of the last probe request and the start of the next request.<br> If the request load is high, the observed delay between requests might significantly exceed the specified interval.<br>3. Make sure the Windows firewall will allow client probing by adding a a remote administration exception to the Windows firewall for each probed client. |
| Step 5 | Save the configuration. | 1. Click **OK** to save the User-ID agent setup settings.<br>2. Click **Commit** to save the configuration. |
| Step 6 | (Optional) Define the set of users for which you do not need to provide IP address to user name mappings, such as kiosk accounts.<br> You can also use the `ignore-user` list to identify users whom you want to force to authenticate using Captive Portal. | 1. Open a CLI session to the firewall.<br>2. To add the list of user accounts for which you do not want the firewall to perform mapping, run the following command:<br>**set user-id-collector ignore-user <value>**<br>Where <value> is a list of the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Separate entries with a space and do not include the domain name with the username. For example:<br>`set user-id-collector ignore-user SPAdmin SPInstall TFSReport`<br>3. **Commit** your changes. |
| Step 7 | Verify the configuration. | 1. From the CLI, enter the following command:<br>`show user server-monitor state all`<br>2. On the **Device > User Identification > User Mapping** tab in the web interface, verify that the **Status** of each server you configured for server monitoring is **Connected**. |

 © Palo Alto Networks, Inc.

# Configure User-ID to Receive User Mappings from a Syslog Sender

The following topics describe how to configure the User-ID agent (either the Windows agent or the integrated agent on the firewall) as a Syslog listener:

- Configure the Integrated User-ID Agent as a Syslog Listener
- Configure the Windows User-ID Agent as a Syslog Listener

## Configure the Integrated User-ID Agent as a Syslog Listener

The following workflow describes how to configure the PAN-OS integrated User-ID agent to receive syslog messages from authenticating services.

The PAN-OS integrated User-ID agent accepts syslogs over SSL and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages. However, if you must use UDP, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending UDP traffic to the firewall.

| **Collect User Mappings from Syslog Senders** | |
|---|---|
| Step 1 Determine whether there is a pre-defined syslog filter for your particular syslog sender(s).<br><br>Palo Alto Networks provides several pre-defined syslog filters, which are delivered as Application content updates and are therefore updated dynamically as new filters are developed. The pre-defined filters are global to the firewall, whereas manually defined filters apply to a single virtual system only.<br><br> Any new syslog filters in a given content release will be documented in the corresponding release note along with the specific regex used to define the filter. | 1. Verify that your Application or Application and Threat database is up to date:<br>a. Select **Device > Dynamic Updates**.<br>b. Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates.<br>c. If a new update is available, **Download** and **Install** it.<br>2. Check to see what pre-defined filters are available:<br>a. Select **Device > User Identification > User Mapping**.<br>b. In the Server Monitoring section of the screen, click **Add**.<br>c. Select **Syslog Sender** as the server **Type**.<br>d. Select the **Filter** drop-down and check to see if there is a filter for the manufacturer and product you plan to forward syslogs from. If the filter you need is available, skip to Step 5 for instructions on defining the servers. If the filter you need is not available, continue to Step 2. |

**Collect User Mappings from Syslog Senders (Continued)**

Step 2 Manually define syslog filter(s) for extracting the User-ID IP address to username mapping information from syslog messages.

In order to be parsed by the User-ID agent, syslog messages must meet the following criteria:

- Each syslog message must be a single line text string. Line breaks are delimited by a carriage return and a new line (`\r\n`) or a new line (`\n`).
- The maximum allowed size of an individual syslog message is 2048 bytes.
- Syslog messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets.
- A single packet may contain multiple syslog messages.

1. Review the syslogs generated by the authenticating service to identify the syntax of the login events. This enables you to create the matching patterns that will allow the firewall to identify and extract the authentication events from the syslogs.

   While reviewing the syslogs also determine whether the domain name is included in the log entries. If the authentication logs do not contain domain information, consider defining a default domain name when adding the syslog sender to the monitored servers list in Step 5.
2. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User-ID Agent Setup section.
3. On the **Syslog Filters** tab, **Add** a new syslog parse profile.
4. Enter a name for the **Syslog Parse Profile**.
5. Specify the **Type** of parsing to use to filter out the user mapping information by selecting one of the following options:
   - **Regex Identifier**—With this type of parsing, you specify regular expressions to describe search patterns for identifying and extracting user mapping information from syslog messages. Continue to Step 3 for instructions on creating the regex identifiers.
   - **Field Identifier**—With this type of parsing, you specify a string to match the authentication event, and prefix and suffix strings to identify the user mapping information in the syslogs. Continue to Step 4 for instructions on creating the field identifiers.

 © Palo Alto Networks, Inc.

**Collect User Mappings from Syslog Senders (Continued)**

Step 3 If you selected **Regex Identifier** as the parsing **Type**, create the regex matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a regex configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event Regex** field. For example, when matched against the example syslog message, the following regex instructs the firewall to extract the first `{1}` instance of the string `authentication success`. The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character: `(authentication\ success){1}`.
2. Enter the regex for identifying the beginning of the username in the authentication success messages in the **Username Regex** field. For example, the regex `User:([a-zA-Z0-9\\\._]+)` would match the string `User:johndoe1` in the example message and extract `acme\johndoe1` as the User-ID.

   If the syslogs do not contain domain information and you require domain names in your user mappings, be sure to enter the **Default Domain Name** when defining the monitored server entry in Step 5.
3. Enter the regex for identifying the IP address portion of the authentication success messages in the **Address Regex** field. For example, the following regular expression `Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})` would match an IPv4 address (`Source:192.168.0.212` in the example syslog).
4. Click **OK**.

**Collect User Mappings from Syslog Senders (Continued)**

Step 4 If you selected **Field Identifier** as the parsing **Type**, define the string matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a field identifier configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event String** field. For example, when matched against the sample syslog message, you would enter the string `authentication success` to identify authentication events in the syslog.
2. Enter the matching string for identifying the beginning of the username field within the authentication syslog message in the **Username Prefix** field. For example, the string `User:` identifies the beginning of the username field in the sample syslog.
3. Enter the **Username Delimiter** to mark the end of the username field within an authentication syslog message. For example, if the username is followed by a space, you would enter `\s` to indicate that the username field is delimited by a standalone space in the sample log.
4. Enter the matching string for identifying the beginning of the IP address field within the authentication event log in the **Address Prefix** field. For example, the string `Source:` identifies the beginning of the address field in the example log.
5. Enter the **Address Delimiter** to mark the end of the IP address field within the authentication success message within the field. For example, if the address is followed by a line break, you would enter `\n` to indicate that the address field is delimited by a new line.
6. Click **OK**.

Step 5 Define the servers that will send syslog messages to the firewall for user mapping purposes.

You can define up to 50 syslog senders per virtual system and up to a total of 100 monitored servers, including syslog senders, Microsoft Active Directory, Microsoft Exchange, or Novell eDirectory servers. The firewall will discard any syslog messages received from servers that are not on this list.

A Syslog sender using SSL to connect will only show a **Status** of **Connected** when there is an active SSL connection. Syslog senders using UDP will not show a **Status** value.

1. Select **Device > User Identification > User Mapping**.
2. In the Server Monitoring section of the screen, click **Add**.
3. Enter a **Name** and **Network Address** for the server.
4. Select **Syslog Sender** as the server **Type**.
5. Make sure the **Enabled** check box is selected.
6. (Optional) If the syslogs that the authenticating device sends do not include domain information in the login event logs, enter the **Default Domain Name** to append to the user mappings.
7. Click **OK** to save the settings.

 © Palo Alto Networks, Inc.

**Collect User Mappings from Syslog Senders (Continued)**

| | |
|---|---|
| Step 6 Enable syslog listener services in the management profile associated with the interface used for user mapping.<br><br>Use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages when using agentless User Mapping on a firewall. However, if you must use UDP, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending UDP traffic to the firewall.<br><br> 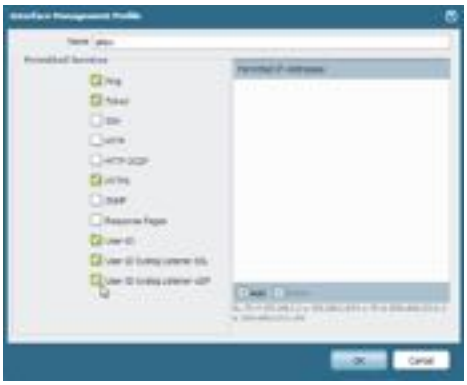 | 1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.<br>2. Select **User-ID Syslog Listener-SSL** and/or **User-ID Syslog Listener-UDP**, depending on the protocols you defined when you set up your Syslog Senders in the Server Monitor list.<br><br>On the Windows User-ID agent, the default listening port for syslog over UDP or TCP is 514, but the port value is configurable. For the agentless User Mapping feature on the firewall only syslog over UDP and SSL are supported and the listening ports (514 for UDP and 6514 for SSL) are not configurable; they are enabled through the management service only.<br><br>3. Click **OK** to save the interface management profile.<br><br>Even after enabling the User-ID Syslog Listener service on the interface, the interface will only accept syslog connections from servers that have a corresponding entry in the User-ID monitored servers configuration. Connections or messages from servers that are not on the list will be discarded. |
| Step 7 Save the configuration. | Click **Commit** to save the configuration. |

**Collect User Mappings from Syslog Senders (Continued)**

Step 8 Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
        UDP Syslog Listener Service is enabled
        SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)     Host: Syslog2(10.5.204.41)
        number of log messages                          : 1000
        number of auth. success messages                : 1000
        number of active connections                    : 0
        total connections made                          : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name                            TYPE     Host            Vsys    Status

------------------------------------------------------------------------------
AD                              AD       10.2.204.43     vsys1   Connected


Syslog Servers:
Name                        Connection Host            Vsys    Status
------------------------------------------------------------------------------
Syslog1                     UDP        10.5.204.40     vsys1   N/A
Syslog2                     SSL        10.5.204.41     vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP              Vsys   From    User                            IdleTimeout(s) M
axTimeout(s)
--------------- ------ ------- ------------------------------- -------------- -
192.168.3.8     vsys1  SYSLOG  acme\jreddick                   2476           2
476
192.168.5.39    vsys1  SYSLOG  acme\jdonaldson                 2480           2
480
192.168.2.147   vsys1  SYSLOG  acme\ccrisp                     2476           2
476
192.168.2.175   vsys1  SYSLOG  acme\jjaso                      2476           2
476
192.168.4.196   vsys1  SYSLOG  acme\jblevins                   2480           2
480
192.168.4.103   vsys1  SYSLOG  acme\bmoss                      2480           2
480
192.168.2.193   vsys1  SYSLOG  acme\esogard                    2476           2
476
192.168.2.119   vsys1  SYSLOG  acme\acallaspo                  2476           2
476
192.168.3.176   vsys1  SYSLOG  acme\jlowrie                    2478           2
478

Total: 9 users
```

 © Palo Alto Networks, Inc.

## Configure the Windows User-ID Agent as a Syslog Listener

The following workflow describes how to configure a Windows-based User-ID agent to listen for syslogs from authenticating services.

The Windows User-ID agent accepts syslogs over TCP and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, use TCP instead of UDP. In either case, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

| **Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders** | |
|---|---|
| Step 1 Launch the User-ID Agent application.<br> | 1. Click Start and select **User-ID Agent**. |
| Step 2 Manually define syslog filter(s) for extracting the User-ID IP address to username mapping information from syslog messages.<br>In order to be parsed by the User-ID agent, syslog messages must meet the following criteria:<br>• Each syslog message must be a single line text string. Line breaks are delimited by a carriage return and a new line (`\r\n`) or a new line (`\n`).<br>• The maximum allowed size of an individual syslog message is 2048 bytes.<br>• Syslog messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets.<br>• A single packet may contain multiple syslog messages. | 1. Review the syslogs generated by the authenticating service to identify the syntax of the login events. This enables you to create the matching patterns that will allow the firewall to identify and extract the authentication events from the syslogs.<br>While reviewing the syslogs also determine whether the domain name is included in the log entries. If the authentication logs do not contain domain information, consider defining a default domain name when adding the syslog sender to the monitored servers list in Step 5.<br>2. Select **User Identification > Setup** and click **Edit** in the Setup section of the dialog.<br>3. On the **Syslog** tab, **Add** a new syslog parse profile.<br>4. Enter a **Profile Name** and **Description**.<br>5. Specify the **Type** of parsing to use to filter out the user mapping information by selecting one of the following options:<br>• **Regex**—With this type of parsing, you specify regular expressions to describe search patterns for identifying and extracting user mapping information from syslog messages. Continue to Step 3 for instructions on creating the regex identifiers.<br>• **Field**—With this type of parsing, you specify a sting to match the authentication event, and prefix and suffix strings to identify the user mapping information in the syslogs. Continue to Step 4 for instructions on creating the field identifiers. |

**Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)**

Step 3 If you selected **Regex** as the parsing **Type**, create the regex matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a regex configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event Regex** field. For example, when matched against the example syslog message, the following regex instructs the firewall to extract the first `{1}` instance of the string `authentication success`. The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character: `(authentication\ success){1}`.
2. Enter the regex for identifying the beginning of the username in the authentication success messages in the **Username Regex** field. For example, the regex `User:([a-zA-Z0-9\\\._]+)` would match the string `User:johndoe1` in the example message and extract `acme\johndoe1` as the User-ID.

   If the syslogs do not contain domain information and you require domain names in your user mappings, be sure to enter the **Default Domain Name** when defining the monitored server entry in Step 5.
3. Enter the regex for identifying the IP address portion of the authentication success messages in the **Address Regex** field. For example, the following regular expression `Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})` would match an IPv4 address (`Source:192.168.0.212` in the example syslog).
4. Click **OK** to save the profile.

 © Palo Alto Networks, Inc.

**Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)**

**Step 4** If you selected **Field Identifier** as the parsing **Type**, define the string matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a field identifier configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event String** field. For example, when matched against the sample syslog message, you would enter the string `authentication success` to identify authentication events in the syslog.
2. Enter the matching string for identifying the beginning of the username field within the authentication syslog message in the **Username Prefix** field. For example, the string `User:` identifies the beginning of the username field in the sample syslog.
3. Enter the **Username Delimiter** to mark the end of the username field within an authentication syslog message. For example, if the username is followed by a space, you would enter `\s` to indicate that the username field is delimited by a standalone space in the sample log.
4. Enter the matching string for identifying the beginning of the IP address field within the authentication event log in the **Address Prefix** field. For example, the string `Source:` identifies the beginning of the address field in the example log.
5. Enter the **Address Delimiter** to mark the end of the IP address field within the authentication success message within the field. For example, if the address is followed by a line break, you would enter `\n` to indicate that the address field is delimited by a new line.
6. Click **OK** to save the profile.

**Step 5** Enable the syslog listening service on the agent.

As a best practice, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

1. Select the **Enable Syslog Service** check box.
2. (Optional) Modify the **Syslog Service Port** number to match the port number used by the syslog sender (default=514).
3. To save the agent syslog configuration, click **OK**.

**Step 6** Define the servers that will send syslog messages to the User-ID agent.

You can define up to 100 syslog senders. The User-ID agent will discard any syslog messages received from servers that are not on this list.

1. Select **User Identification > Discovery**.
2. In the **Servers** section of the screen, click **Add**.
3. Enter a **Name** and **Server Address** for the server that will send syslogs to the agent.
4. Select **Syslog Sender** as the **Server Type**.
5. Select a **Filter** you defined in Step 2.
6. (Optional) If the syslogs that the authenticating device sends do not include domain information in the login event logs, enter the **Default Domain Name** to append to the user mappings.
7. Click **OK** to save the settings.

**Step 7** Save the configuration.

Click **Commit** to save the configuration.

**Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)**

Step 8 Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
        UDP Syslog Listener Service is enabled
        SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)     Host: Syslog2(10.5.204.41)
        number of log messages                          : 1000
        number of auth. success messages                : 1000
        number of active connections                    : 0
        total connections made                          : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name                          TYPE     Host            Vsys    Status

------------------------------------------------------------------------------
AD                            AD       10.2.204.43     vsys1   Connected


Syslog Servers:
Name                      Connection Host            Vsys    Status
------------------------------------------------------------------------------
Syslog1                   UDP        10.5.204.40     vsys1   N/A
Syslog2                   SSL        10.5.204.41     vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP              Vsys   From    User                            IdleTimeout(s) M
axTimeout(s)
--------------- ------ ------- ------------------------------- -------------- -
192.168.3.8     vsys1  SYSLOG  acme\jreddick                   2476           2
476
192.168.5.39    vsys1  SYSLOG  acme\jdonaldson                 2480           2
480
192.168.2.147   vsys1  SYSLOG  acme\ccrisp                     2476           2
476
192.168.2.175   vsys1  SYSLOG  acme\jjaso                      2476           2
476
192.168.4.196   vsys1  SYSLOG  acme\jblevins                   2480           2
480
192.168.4.103   vsys1  SYSLOG  acme\bmoss                      2480           2
480
192.168.2.193   vsys1  SYSLOG  acme\esogard                    2476           2
476
192.168.2.119   vsys1  SYSLOG  acme\acallaspo                  2476           2
476
192.168.3.176   vsys1  SYSLOG  acme\jlowrie                    2478           2
478

Total: 9 users
```

 © Palo Alto Networks, Inc.

# Map IP Addresses to User Names Using Captive Portal

If the firewall receives a request from a zone that has User-ID enabled and the source IP address does not have any user data associated with it yet, it checks its Captive Portal policy for a match to determine whether to perform authentication. This is useful in environments where you have clients that are not logged in to your domain servers, such as Linux clients. This user mapping method is only triggered for web traffic (HTTP or HTTPS) that matches a security rule/policy, but that has not been mapped using a different method.

- Captive Portal Authentication Methods
- Captive Portal Modes
- Configure Captive Portal

## Captive Portal Authentication Methods

Captive Portal uses the following methods to obtain user data from the client when a request matches a Captive Portal policy:

| Authentication Method | Description |
|---|---|
| NT LAN Manager (NTLM) | The firewall uses an encrypted challenge-response mechanism to obtain the user's credentials from the browser. When configured properly, the browser will provide the credentials to the firewall transparently without prompting the user, but will display a prompt for credentials if necessary. If the browser cannot perform NTLM or if NTLM authentication fails, the firewall falls back to web form or client certificate authentication, depending on your Captive Portal configuration.<br>By default, IE supports NTLM. Firefox and Chrome can be configured to use it. You cannot use NTLM to authenticate non-Windows clients. |
| Web Form | Requests are redirected to a web form for authentication. You can configure Captive Portal to use a local user database, RADIUS, LDAP, or Kerberos to authenticate users. Although users will always be prompted for credentials, this authentication method works with all browsers and operating systems. |
| Client Certificate Authentication | Prompts the browser to present a valid client certificate for authenticating the user. To use this method, you must provision client certificates on each user system and install the trusted CA certificate used to issue those certificates on the firewall. This is the only authentication method that enables transparent authentication for Mac OS and Linux clients. |

## Captive Portal Modes

The Captive Portal mode defines how web requests are captured for authentication:

© Palo Alto Networks, Inc.

| Mode | Description |
|---|---|
| Transparent | The firewall intercepts the browser traffic per the Captive Portal rule and impersonates the original destination URL, issuing an HTTP 401 to invoke authentication. However, because the firewall does not have the real certificate for the destination URL, the browser will display a certificate error to users attempting to access a secure site. Therefore you should only use this mode when absolutely necessary, such as in Layer 2 or virtual wire deployments. |
| Redirect | The firewall intercepts unknown HTTP or HTTPS sessions and redirects them to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to perform authentication. This is the preferred mode because it provides a better end-user experience (no certificate errors). However, it does require additional Layer 3 configuration. Another benefit of the Redirect mode is that it provides for the use of session cookies, which enable the user to continue browsing to authenticated sites without requiring re-mapping each time the time outs expire. This is especially useful for users who roam from one IP address to another (for example, from the corporate LAN to the wireless network) because they will not need to re-authenticate upon IP address change as long as the session stays open. In addition, if you plan to use NTLM authentication, you must use Redirect mode because the browser will only provide credentials to trusted sites. |

## Configure Captive Portal

The following procedure shows how to configure Captive Portal using the PAN-OS integrated User-ID agent to redirect requests that match a Captive Portal policy to a Layer 3 interface on the firewall.

If you plan to use Captive Portal without using the other User-ID functions (user mapping and group mapping), you do not need to configure an agent.

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent | | |
|---|---|---|
| Step 1 | Make sure the firewall has a route to the servers it will be monitoring to gather user data (for example, your Domain Controllers and your Exchange servers). | In this release of the product, the firewall must be able to communicate with the servers over the MGT interface, so you must make sure that the network your directory servers are on is accessible from this interface. If this configuration does not work in your environment, you must configure Captive Portal using the Window-based User-ID agent. |
| Step 2 | Make sure DNS is configured to resolve your Domain Controller addresses. | To verify proper resolution, ping the server FQDN. For example:<br>`admin@PA-200> ping host dc1.acme.com` |

 © Palo Alto Networks, Inc.

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | | |
|---|---|---|
| Step 3 | (Redirect mode only) Create a Layer 3 interface to which to redirect Captive Portal requests. | 1. Create a management profile to enable the interface to display Captive Portal response pages:<br>a. Select **Network > Interface Mgmt** and click **Add**.<br>b. Enter a **Name** for the profile, select **Response Pages**, and then click **OK**.<br>2. Create the Layer 3 interface. Be sure to attach the management profile you just created (on the **Advanced > Other Info** tab of the Ethernet Interface dialog).<br>3. Create a DNS "A" record that maps the IP address you configured on the Layer 3 interface to an intranet host name (that is, a hostname that does not have a dot in the name, such as `ntlmhost`). |
| Step 4 | (Redirect mode only) To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface to which you are redirecting requests.You can either generate a self-signed certificate or import a certificate that is signed by an external CA.<br>When setting up Captive Portal for the first time, imported certificates may not work. If you plan to use an imported certificate, complete the initial configuration without specifying a **Server Certificate**. After you get Captive Portal working, you can go back and switch to the imported certificate. | To use a self-signed certificate, you must first create a root CA certificate and then use that CA to sign the certificate you will use for Captive Portal as follows:<br>1. To create a root CA certificate, select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**. Enter a **Certificate Name**, such as RootCA. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed). Make sure you select the **Certificate Authority** check box and then click **Generate** the certificate.<br>2. To create the certificate to use for Captive Portal, click **Generate**. Enter a **Certificate Name** and enter the DNS name of the intranet host for the interface as the **Common Name**. In the **Signed By** field, select the CA you created in the previous step. Add an IP address attribute and specify the IP address of the Layer 3 interface to which you will be redirecting requests. **Generate** the certificate.<br> 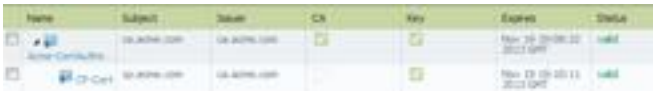 <br>3. To configure clients to trust the certificate, select the CA certificate on the **Device Certificates** tab and click **Export**. You must then import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO). |

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | | |
|---|---|---|
| Step 5 | Set up an authentication mechanism to use when the web form is invoked. Note that even if you plan to use NTLM, you must also set up a secondary authentication mechanism that can be used if NTLM authentication fails or if the user agent does not support it.<br><br>**Best Practices:**<br>•If using RADIUS to authenticate users from the web form, be sure to enter a RADIUS domain. This will be used as the default domain if users don't supply one upon login.<br>•If using AD to authenticate users from the web form, make sure to enter **sAMAccountName** as the **LogonAttribute**. | 1. Configure the firewall to connect to the authentication service you plan to use so that it can access the authentication credentials.<br>• If you plan to authenticate using LDAP, Kerberos, or RADIUS you must create a server profile that instructs the firewall how to connect to the service and access the authentication credentials for your users. Select **Device > Server Profiles** and add a new profile for the specific service you will be accessing.<br>• If you plan to use local database authentication, you must first create the local database. Select **Device > Local User Database** and add the users and groups to be authenticated.<br>2. Create an authentication profile that references the server profile or local user database you just created. Select **Device > Authentication Profile** and add a new profile for use with Captive Portal. For details on creating a specific type of authentication profile, refer to the online help. |
| Step 6 | (Optional) Set up client certificate authentication. Note that you do not need to set up both an authentication profile and a client certificate profile to enable Captive Portal. If you configure both, the user will be required to authenticate using both methods.<br><br>For details on other certificate profile fields, such as whether to use CRL or OCSP, refer to the online help. | 1. Generate certificates for each user who will be authenticating using Captive Portal.<br>2. Download the CA certificate in Base64 format.<br>3. Import the root CA certificate from the CA that generated the client certificates onto the firewall:<br>a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.<br>c. **Browse** to the **Certificate File** you downloaded from the CA.<br>d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.<br>e. Select the certificate you just imported on the **Device Certificates** tab to open it.<br>f. Select **Trusted Root CA** and then click **OK**.<br>4. Create the client certificate profile that you will use when you configure Captive Portal.<br>a. Select **Device > Certificates > Certificate Management > Certificate Profile** and click **Add** and enter a profile **Name**.<br>b. In the **Username Field** drop-down, select the certificate field that contains the user's identity information.<br>c. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you just imported and then click **OK**. |

 © Palo Alto Networks, Inc.

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | |
|---|---|
| Step 7 Enable NTLM authentication.<br> When using the on-device User-ID agent, the firewall must be able to successfully resolve the DNS name of your Domain Controller in order for the firewall to join the domain. The credentials you supply here will be used to join the firewall to the domain upon successful DNS resolution. | 1. Select **Device > User Identification > User Mapping** and click the Edit icon in the **Palo Alto Networks User ID Agent Setup** section of the screen.<br>2. On the **NTLM** tab, select the **Enable NTLM authentication processing** check box.<br>3. Enter the NTLM domain against which the User-ID agent on the firewall should check NTLM credentials.<br>4. Enter the user name and password for the Active Directory account you created in Step 1 in Map IP Addresses to Users Using the Integrated User-ID Agent for NTLM authentication.<br> Do not include the domain in the **Admin User Name** field. Otherwise, the firewall will fail to join the domain. |

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | |
|---|---|
| Step 8 Configure the Captive Portal settings. | 1. Select **Device > User Identification > Captive Portal Settings** and click the Edit icon in the **Captive Portal** section of the screen.<br>2. Make sure the **Enabled** check box is selected.<br>3. Set the **Mode**. This example shows how to set up **Redirect** mode.<br>4. (Redirect mode only) Select the **Server Certificate** the firewall should use to redirect requests over SSL. This is the certificate you created in Step 4.<br>5. (Redirect mode only) Specify the **Redirect Host**, which is the intranet hostname that resolves to the IP address of the Layer 3 interface to which you are redirecting requests, as specified in Step 3.<br>6. Select the authentication method to use if NTLM fails (or if you are not using NTLM):<br>• If you are using LDAP, Kerberos, RADIUS, or local database authentication, select the **Authentication Profile** you created in Step 5.<br>• If you are using client certificate authentication, select the **Certificate Profile** you created in Step 6.<br> 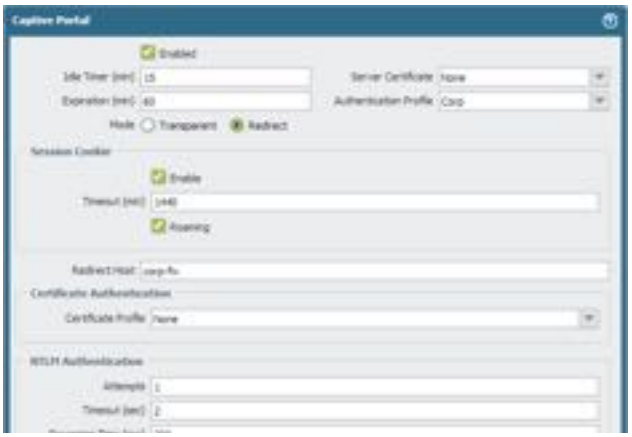 <br>7. Click **OK** to save your settings.<br>8. Click **Commit** to save the Captive Portal configuration. |

 © Palo Alto Networks, Inc.

# Configure User Mapping for Terminal Server Users

Individual terminal server users appear to have the same IP address and therefore an IP address to username mapping is not sufficient to identify a specific user. To enable identification of specific users on Windows-based terminal servers, the Palo Alto Networks Terminal Services agent (TS agent) allocates a port range to each user. It then notifies every connected firewall about the allocated port range, which allows the firewall to create an IP address-port-user mapping table and enable user- and group-based security policy enforcement. For non-Windows terminal servers you can configure the User-ID XML API to extract user mapping information.

The following sections describe how to configure user mapping for terminal server users:

- Configure the Palo Alto Networks Terminal Server Agent for User Mapping
- Retrieve User Mappings from a Terminal Server Using the User-ID XML API

## Configure the Palo Alto Networks Terminal Server Agent for User Mapping

Use the following procedure to install the TS agent on the terminal server. You must install the TS agent on all terminal servers that your users log in to in order to successfully map all your users.

For information about the supported terminal servers supported by the TS Agent, refer to "Operating System (OS) Compatibility TS Agent" in the Terminal Services Agent Release Notes, which are available on the Palo Alto Networks Software Updates page.

| **Install the Windows Terminal Server Agent** | |
|---|---|
| Step 1 Download the TS Agent installer. | 1. Log in to the Palo Alto Networks Support site.<br>2. Select **Software Updates** from the Manage Devices section.<br>3. Scroll to the **Terminal Services Agent** section and **Download** the version of the agent you want to install.<br>4. Save the `TaInstall64.x64-x.x.x-xx.msi` or `TaInstall-x.x.x-xx.msi` file (be sure to select the appropriate version based on whether the Windows system is running a 32-bit OS or a 64-bit OS) on the system(s) where you plan to install the agent. |

| Install the Windows Terminal Server Agent (Continued) | |
|---|---|
| Step 2 Run the installer as an administrator.<br> | 1. To launch a command prompt as an administrator, click Start and right-click **Command Prompt** and then select **Run as administrator**.<br>2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop you would enter the following:<br>`C:\Users\administrator.acme>cd Desktop`<br>`C:\Users\administrator.acme\Desktop>TaInstall-6.0.0-1.msi`<br>3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to the `C:\Program Files (x86)\Palo Alto Networks\Terminal Server Agent` folder, but you can **Browse** to a different location.<br>4. When the installation completes, **Close** the setup window.<br>If you are upgrading to a TS Agent version that has a newer driver than the existing installation, the installation wizard prompts you to reboot the system after upgrading in order to use the new driver. |
| Step 3 Launch the Terminal Server Agent application. | Click Start and select **Terminal Server Agent**.<br> 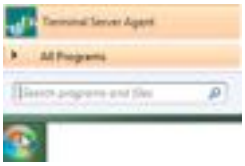 |

 © Palo Alto Networks, Inc.

**Install the Windows Terminal Server Agent (Continued)**

Step 4 Define the range of ports for the TS Agent to allocate to end users.

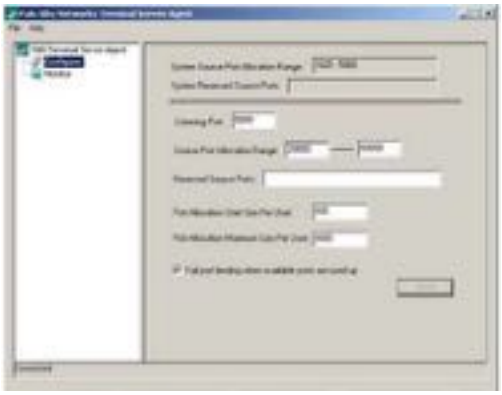

The **System Source Port Allocation Range** and **System Reserved Source Ports** fields specify the range of ports that will be allocated to non-user sessions. Make sure the values specified in these fields do not overlap with the ports you designate for user traffic. These values can only be changed by editing the corresponding Windows registry settings.

1. Select **Configure**.
2. Set the **Source Port Allocation Range** (default 20000-39999). This is the full range of port numbers that the TS Agent will allocate for user mapping. The port range you specify cannot overlap with the **System Source Port Allocation Range**.
3. (Optional) If there are ports/port ranges within the source port allocation that you do not want the TS Agent to allocate to user sessions, specify them as **Reserved Source Ports**. To include multiple ranges, use commas with no spaces, for example: `2000-3000,3500,4000-5000`.
4. Specify the number of ports to allocate to each individual user upon login to the terminal server in the **Port Allocation Start Size Per User** field (default 200).
5. Specify the **Port Allocation Maximum Size Per User**, which is the maximum number of ports the Terminal Server agent can allocate to an individual user.
6. Specify whether to continue processing traffic from the user if the user runs out of allocated ports. By default, the **Fail port binding when available ports are used up** is selected, which indicates that the application will fail to send traffic when all ports are used. To enable users to continue using applications when they run out of ports, clear this check box. Keep in mind that this traffic may not be identified with User-ID.

Step 5 Configure the firewalls to connect to the Terminal Server agent.

Complete the following steps on each firewall you want to connect to the Terminal Server agent to receive user mappings:

1. Select **Device > User Identification > Terminal Server Agents** and click **Add**.
2. Enter a **Name** for the Terminal Server agent.
3. Enter the IP address of the Windows **Host** on which the Terminal Server agent is installed.
4. Enter the **Port** number on which the agent will listen for user mapping requests. This value must match the value configured on the Terminal Server agent. By default, the port is set to 5009 on the firewall and on the agent. If you change it here you must also change the **Listening Port** field on the Terminal Server agent **Configure** screen.
5. Make sure that the configuration is **Enabled** and then click **OK**.
6. **Commit** the changes.
7. Verify that the **Connected** status displays as connected.

| Install the Windows Terminal Server Agent (Continued) | |
|---|---|
| Step 6 Verify that the Terminal Server agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.<br> 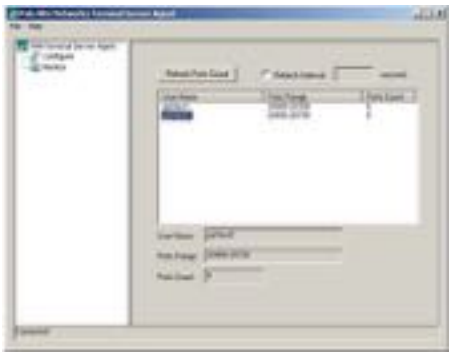 | 1. Launch the Terminal Server agent and verify that the firewalls can connect by making sure the **Connection Status** for each of Device in the Connection List is **Connected**.<br>2. To verify that the Terminal Server agent is successfully mapping port ranges to usernames, select **Monitoring** and make sure that the mapping table is populated. |

## Retrieve User Mappings from a Terminal Server Using the User-ID XML API

The User-ID XML API is a RESTful API that uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services.

To enable a non-Windows terminal server to send user mapping information directly to the firewall, create scripts that extract the user login and logout events and use them for input to the User-ID XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget and providing the firewall's API key for secure communication. Creating user mappings from multi-user systems such as terminal servers requires use of the following API messages:

- **<multiusersystem>**—Sets up the configuration for an XML API Multi-user System on the firewall. This message allows for definition of the terminal server IP address (this will be the source address for all users on that terminal server). In addition, the <multiusersystem> setup message specifies the range of source port numbers to allocate for user mapping and the number of ports to allocate to each individual user upon login (called the *block size*). If you want to use the default source port allocation range (1025-65534) and block size (200), you do not need to send a <multiusersystem> setup event to the firewall. Instead, the firewall will automatically generate the XML API Multi-user System configuration with the default settings upon receipt of the first user login event message.
- **<blockstart>**—Used with the <login> and <logout> messages to indicate the starting source port number allocated to the user. The firewall then uses the block size to determine the actual range of port numbers to map to the IP address and username in the login message. For example, if the <blockstart> value is 13200 and the block size configured for the multi-user system is 300, the actual source port range allocated to the user is 13200 through 13499. Each connection initiated by the user should use a unique source port number within the allocated range, enabling the firewall to identify the user based on its IP address-port-user mappings for enforcement of user- and group-based security policy rules. When a user exhausts all the ports allocated, the terminal server must send a new <login> message allocating a new port range for the user so that the firewall can update the IP address-port-user mapping. In addition, a single

 © Palo Alto Networks, Inc.

username can have multiple blocks of ports mapped simultaneously. When the firewall receives a <logout> message that includes a <blockstart> parameter, it removes the corresponding IP address-port-user mapping from its mapping table. When the firewall receives a <logout> message with a username and IP address, but no <blockstart>, it removes the user from its table. And, if the firewall receives a <logout> message with an IP address only, it removes the multi-user system and all mappings associated with it.

The XML files that the terminal server sends to the firewall can contain multiple message types and the messages do not need to be in any particular order within the file. However, upon receiving an XML file that contains multiple message types, the firewall will process them in the following order: multiusersystem requests first followed by logins then logouts.

The following workflow provides an example of how to use the User-ID XML API to send user mappings from a non-Windows terminal server to the firewall.

**Use the User-ID XML API to Map Non-Windows Terminal Services Users**

Step 1 — Generate the API key that will be used to authenticate the API communication between the firewall and the Terminal server. To generate the key you must provide login credentials for an administrative account; the API is available to all administrators (including role-based administrators with XML API privileges enabled).

Any special characters in the password must be URL/percent-encoded.

From a browser, log in to the firewall. Then, to generate the API key for the firewall, open a new browser window and enter the following URL:

```
https://<Firewall-IPaddress>/api/?type=keygen&user=<username>&
password=<password>
```

Where <Firewall-IPaddress> is the IP address or FQDN of the firewall and <username> and <password> are the credentials for the administrative user account on the firewall. For example:

```
https://10.1.2.5/api/?type=keygen&user=admin&password=admin
```

The firewall responds with a message containing the key, for example:

```
<response status="success">
   <result>
      <key>k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg=</key>
   </result>
</response>
```

**Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)**

Step 2 (Optional) Generate a setup message that the terminal server will send to specify the port range and block size of ports per user that your terminal services agent uses.

If the terminal services agent does not send a setup message, the firewall will automatically create a terminal server agent configuration using the following default settings upon receipt of the first login message:

- Default port range: 1025 to 65534
- Per user block size: 200
- Maximum number of multi-user systems: 1000

The following shows a sample setup message:

```
<uid-message>
  <payload>
    <multiusersystem>
      <entry ip="10.1.1.23" startport="20000"
        endport="39999" blocksize="100">
    </multiusersystem>
  </payload>
  <type>update</type>
  <version>1.0</version>
</uid-message>
```

where `entry ip` specifies the IP address assigned to terminal server users, `startport` and `endport` specify the port range to use when assigning ports to individual users and `blocksize` specifies the number of ports to assign to each user. The maximum blocksize is 4000 and each multi-user system can allocate a maximum of 1000 blocks.

If you define a custom blocksize and or port range, keep in mind that you must configure the values such that every port in the range gets allocated and that there are no gaps or unused ports. For example, if you set the port range to 1000-1499, you could set the block size to 100, but not to 200. This is because if you set it to 200, there would be unused ports at the end of the range.

Step 3 Create a script that will extract the login events and create the XML input file to send to the firewall.

Make sure the script enforces assignment of port number ranges at fixed boundaries with no port overlaps. For example, if the port range is 1000-1999 and the block size is 200, acceptable blockstart values would be 1000, 1200, 1400, 1600, or 1800. Blockstart values of 1001, 1300, or 1850 would be unacceptable because some of the port numbers in the range would be left unused.

The login event payload that the terminal server sends to the firewall can contain multiple login events.

The following shows the input file format for a user-ID XML login event:

```
<uid-message>
<payload>
<login>
<entry name="acme\jjaso" ip="10.1.1.23" blockstart="20000">
<entry name="acme\jparker" ip="10.1.1.23" blockstart="20100">
<entry name="acme\ccrisp" ip="10.1.1.23" blockstart="21000">
</login>
</payload>
<type>update</type>
<version>1.0</version>
</uid-message>
```

The firewall uses this information to populate its user mapping table. Based on the mappings extracted from the example above, if the firewall received a packet with a source address and port of 10.1.1.23:20101, it would map the request to user jparker for policy enforcement.

Each multi-user system can allocate a maximum of 1000 port blocks.

 © Palo Alto Networks, Inc.

**Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)**

| | | |
|---|---|---|
| Step 4 | Create a script that will extract the logout events and create the XML input file to send to the firewall.<br><br>Upon receipt of a `logout` event message with a `blockstart` parameter, the firewall removes the corresponding IP address-port-user mapping. If the `logout` message contains a username and IP address, but no `blockstart` parameter, the firewall removes all mappings for the user. If the `logout` message contains an IP address only, the firewall removes the multi-user system and all associated mappings. | The following shows the input file format for a User-ID XML logout event:<br>`<uid-message>`<br>`  <payload>`<br>`  <logout>`<br>`  <entry name="acme\jjaso" ip="10.1.1.23" blockstart="20000">`<br>`  <entry name="acme\ccrisp" ip="10.1.1.23">`<br>`  <entry ip="10.2.5.4">`<br>`  </logout>`<br>`  </payload>`<br>`  <type>update</type>`<br>`  <version>1.0</version>`<br>`</uid-message>`<br><br>You can also clear the multiuser system entry from the firewall using the following CLI command: **`clear xml-api multiusersystem`** |
| Step 5 | Make sure that the scripts you create include a way to dynamically enforce that the port block range allocated using the XML API matches the actual source port assigned to the user on the terminal server and that the mapping is removed when the user logs out or the port allocation changes. | One way to do this would be to use netfilter NAT rules to hide user sessions behind the specific port ranges allocated via the XML API based on the uid. For example, to ensure that a user with the user ID jjaso is mapped to a source network address translation (SNAT) value of 10.1.1.23:20000-20099 the script you create should include the following:<br>`[root@ts1 ~]# iptables -t nat -A POSTROUTING -m owner --uid-owner jjaso -p tcp -j SNAT --to-source 10.1.1.23:20000-20099`<br>Similarly, the scripts you create should also ensure that the IP table routing configuration dynamically removes the SNAT mapping when the user logs out or the port allocation changes:<br>`[root@ts1 ~]# iptables -t nat -D POSTROUTING 1` |
| Step 6 | Define how to package the XML input files containing the setup, login, and logout events into wget or cURL messages for transmission to the firewall. | **To apply the files to the firewall using wget:**<br>`> wget --post file <filename> "https://<Firewall-IPaddress>/api/?type=user-id&key=<key>&file-name=<input_filename.xml>&client=wget&vsys=<VSYS_name>"`<br>For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key `k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg` using wget would look as follows:<br>`> wget --post file login.xml "https://10.2.5.11/api/?type=user-id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg&file-name=login.xml&client=wget&vsys=vsys1"`<br><br>**To apply the file to the firewall using cURL:**<br>`> curl --form file=@<filename> https://<Firewall-IPaddress>/api/?type=user-id&key=<key>&vsys=<VSYS_name>`<br>For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key `k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg` using cURL would look as follows:<br>`> curl --form file@login.xml "https://10.2.5.11/api/?type=user-id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg&vsys=vsys1"` |

| Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued) | | |
|---|---|---|
| Step 7 | Verify that the firewall is successfully receiving login events from the terminal servers. | Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:<br>**To verify if the terminal server is connecting to the firewall over XML:**<br>`admin@PA-5050> show user xml-api multiusersystem`<br>`Host Vsys Users Blocks`<br>`----------`<br>`10.5.204.43 vsys1 5 2`<br>**To verify that the firewall is receiving mappings from a terminal server over XML:**<br>`admin@PA-5050> show user ip-port-user-mapping all`<br>`Global max host index 1, host hash count 1`<br>`XML API Multi-user System 10.5.204.43`<br>`Vsys 1, Flag 3`<br>`Port range: 20000 - 39999`<br>`Port size: start 200; max 2000`<br>`Block count 100, port count 20000`<br>`20000-20199: acme\administrator`<br>`Total host: 1` |

 © Palo Alto Networks, Inc.

## Send User Mappings to User-ID Using the XML API

Although the User-ID functionality provides many out-of-the box methods for obtaining user mapping information, you may have some applications or devices that capture user information that cannot be natively integrated with User-ID. In this case you can use the User-ID XML API to create custom scripts that allow you to leverage existing user data and send it to the User-ID agent or directly to the firewall.

The User-ID XML API is a RESTful API that uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services. To leverage user data from an existing system—such as a custom application developed internally or another device that is not supported by one of the existing user mapping mechanisms—you can create custom scripts to extract the data and send it to the firewall or the User-ID agent using the XML API.

To enable an external system to send user mapping information to the User-ID agent or directly to the firewall, you can create scripts that extract the user login and logout events and use them for input to the User-ID XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget using the firewall's API key for secure communication. For more details, refer to the PAN-OS XML API Usage Guide.

## Configure a Firewall to Share User Mapping Data with Other Firewalls

Because policy is local to each firewall, each firewall needs current user mapping and group mapping information to accurately enforce security policy by user and group. However, you can configure one firewall to collect all the mapping information and act as a User-ID agent to share that information with other firewalls. The redistribution firewall can share only the information it collects using local methods (for example, the PAN-OS integrated User-ID agent or Captive Portal), not the information collected from the Windows-based User-ID and Terminal Services agents. You configure the receiving firewalls to retrieve the mapping information from the redistribution firewall; they don't need to communicate directly with domain servers.

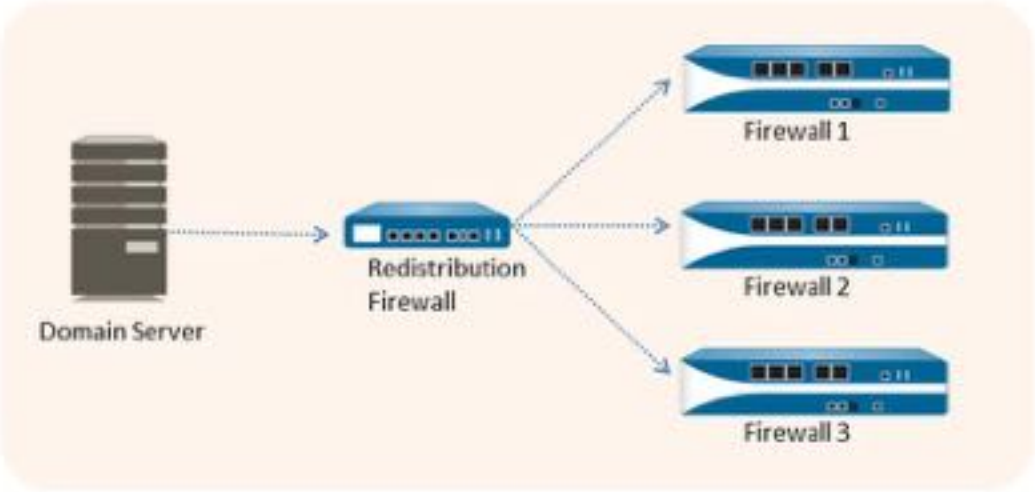


The following procedure describes how to set up redistribution of User-ID information.

| Configure a Firewall to Redistribute User Mappings | |
|---|---|
| Step 1 Configure the redistribution firewall.<br>User-ID configurations apply to a single virtual system only. To redistribute User-ID mappings from multiple virtual systems you must configure the user mapping settings on each virtual system separately, using a unique pre-shared key in each configuration. | 1. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User-ID Agent Setup section.<br>2. Select **Redistribution**.<br>3. Enter a **Collector Name**.<br>4. Enter and confirm the **Pre-Shared Key** that will enable other firewalls to connect to this firewall to retrieve user mapping information.<br>5. Click **OK** to save the redistribution configuration. |

 © Palo Alto Networks, Inc.

| Configure a Firewall to Redistribute User Mappings (Continued) | | |
|---|---|---|
| Step 2 | Create an interface management profile that enables the User-ID service and attach it to the interface that the other firewalls will connect to in order to retrieve user mappings. | 1. Select **Network > Network Profiles > Interface Mgmt** and click **Add**.<br>2. Enter a **Name** for the profile and then select the Permitted Services. At a minimum, select **User-ID Service** and **HTTPS**.<br>3. Click **OK** to save the profile.<br>4. Select **Network > Interfaces > Ethernet** and select the interface you plan to use for redistribution.<br>5. On the **Advanced > Other Info** tab, select the **Management Profile** you just created.<br>6. Click **OK** and **Commit**. |
| Step 3 | Configure the other firewalls to retrieve user mappings from the redistribution firewall.<br> If the redistribution firewall has multiple virtual systems configured for redistribution, make sure you are using the pre-shared key that corresponds to the virtual system from which you want this firewall to retrieve User-ID mappings. | Perform the following steps on each firewall that you want to be able to retrieve user mappings:<br>1. Select **Device > User Identification > User-ID Agents**.<br>2. Click **Add** and enter a User-ID agent **Name** for the redistribution firewall.<br>3. Enter the hostname or IP address of the firewall interface that you configured for redistribution in the **Host** field.<br>4. Enter 5007 as the **Port** number on which the redistribution firewall will listen for User-ID requests.<br>5. Enter the **Collector Name** that you specified in the redistribution firewall configuration (Step 1-3).<br>6. Enter and confirm the **Collector Pre-Shared Key**. The key value you enter here must match the value configured on the redistribution firewall (Step 1-4).<br>7. (Optional) If you are using the redistribution firewall to retrieve group mappings in addition to user mappings, select the **Use as LDAP Proxy** check box.<br>8. (Optional) If you are using the redistribution firewall for Captive Portal authentication, select the **Use for NTLM Authentication** check box.<br>9. Make sure that the configuration is **Enabled** and then click **OK**.<br>10. **Commit** the changes. |
| Step 4 | Verify the configuration. | On the **User-ID Agents** tab, verify that the redistribution firewall entry you just added shows a green icon in the **Connected** column.<br>If a red icon appears, check traffic logs (**Monitor > Logs > Traffic**) to identify the issue. You can also check to see if any user mapping data has been received by running the following operational commands from the CLI:<br>`show user ip-user-mapping` (to view user mapping information on the dataplane)<br>`show user ip-user-mapping-mp` (to view mappings on the management plane). |

# Enable User- and Group-Based Policy

In order to enable security policy based on user and/or group, you must enable User-ID for each zone that contains users you want to identify. You can then define policies that allow or deny traffic based on user name or group membership. Additionally, you can create Captive Portal policies to enable identification for IP addresses that do not yet have any user data associated with them.

| Enable User- and Group-Based Policy | |
|---|---|
| Step 1 Enable User-ID on the source zones that contain the users that will send requests that require user-based access controls. | 1. Select **Network > Zones**.<br>2. Click the **Name** of the zone in which you want to enable User-ID to open the Zone dialog.<br>3. Select the **Enable User Identification** check box and then click **OK**.<br> 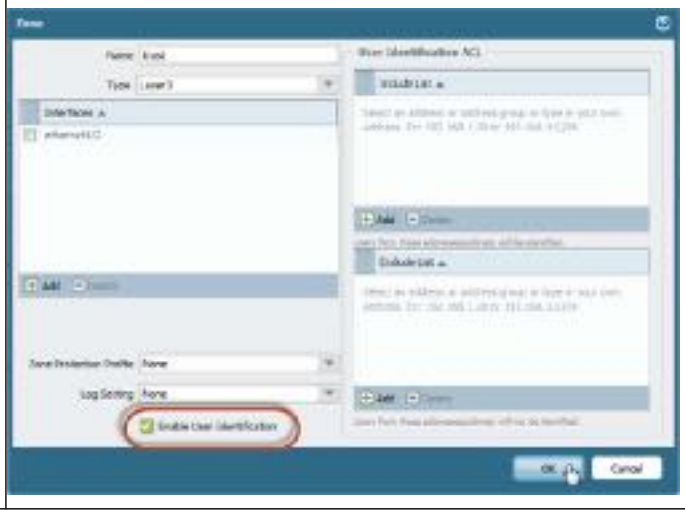 |

 © Palo Alto Networks, Inc.

| Enable User- and Group-Based Policy (Continued) | |
|---|---|
| Step 2 Create security policies based on user and/or group.<br><br> As a best practice, create policies based on group rather than user whenever possible. This prevents you from having to continually update your policies (which requires a commit) whenever your user base changes. | 1. After configuring User-ID, you will be able to choose a user name or group name when defining the source or destination of a security rule:<br>a. Select **Policies > Security** and click **Add** to create a new policy or click on an existing policy rule name to open the Security Policy Rule dialog.<br>b. Specify which users and/or groups to match in the policy in one of the following ways:<br>– If you want to specify specific users/groups as matching criteria, select the **User** tab and click the **Add** button in the Source User section to display a list of users and groups discovered by the firewall group mapping function. Select the users and/or groups to add to the policy.<br>– If you want the policy to match any user who has or has not successfully authenticated and you don't need to know the specific user or group name, select **known-user** or **unknown** from the drop-down list above the **Source User** list.<br>2. Configure the rest of the policy as appropriate and then click **OK** to save it. For details on other fields in the security policy, see Set Up Basic Security Policies. |

| Enable User- and Group-Based Policy (Continued) | |
|---|---|
| Step 3 Create your Captive Portal policies. | 1. Select **Policies > Captive Portal**.<br>2. Click **Add** and enter a **Name** for the policy.<br>3. Define the matching criteria for the rule by completing the **Source**, **Destination**, and **Service/URL Category** tabs as appropriate to match the traffic you want to authenticate. The matching criteria on these tabs is the same as the criteria you define when creating a security policy. See Set Up Basic Security Policies for details.<br>4. Define the **Action** to take on traffic that matches the rule:<br>• **no-captive-portal**—Allow traffic to pass without presenting a Captive Portal page for authentication.<br>• **web-form**—Present a Captive Portal page for the user to explicitly enter authentication credentials or use client certificate authentication.<br>• **browser-challenge**—Open an NTLM authentication request to the web browser of the user. The web browser will respond using the current login credentials of the user. If the login credentials are not available, the user will be prompted to supply them. If you select this action, you must enable NT LAN Manager (NTLM) authentication when you Configure Captive Portal. If NTLM authentication fails, the firewall falls back to **web-form** authentication.<br>5. Click **OK**.<br>The following example shows a Captive Portal policy that instructs the firewall to present a web form to authenticate unknown users who send HTTP requests from the trust zone to the untrust zone.<br> 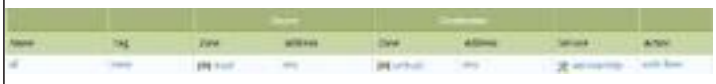 |
| Step 4 Save your policy settings. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# Verify the User-ID Configuration

After you configure group mapping and user mapping and enable User-ID on your security policies and Captive Portal policies, you should verify that it is working properly.

| Verify the User-ID Configuration | |
|---|---|
| Step 1 Verify that group mapping is working. | From the CLI, enter the following command:<br>**show user group-mapping statistics** |
| Step 2 Verify that user mapping is working. | If you are using the on-device User-ID agent, you can verify this from the CLI using the following command:<br>**show user ip-user-mapping-mp all**<br><br>IP Vsys From User Timeout (sec)<br>---<br>192.168.201.1 vsys1 UIA acme\george 210<br>192.168.201.11 vsys1 UIA acme\duane 210<br>192.168.201.50 vsys1 UIA acme\betsy 210<br>192.168.201.10 vsys1 UIA acme\administrator 210<br>192.168.201.100 vsys1 AD acme\administrator 748<br>Total: 5 users<br>*: WMI probe succeeded |

| Verify the User-ID Configuration (Continued) | |
|---|---|
| Step 3 Test your security policy. | • From a machine in the zone where User-ID is enabled, attempt to access sites and applications to test the rules you defined in your policy and ensure that traffic is allowed and denied as expected.<br>• You can also use the `test security-policy-match` command to determine whether the policy is configured correctly. For example, suppose you have a rule that blocks user duane from playing World of Warcraft, you could test the policy as follows:<br>`test security-policy-match application worldofwarcraft source-user acme\duane source any destination any destination-port any protocol 6`<br><br>`"deny worldofwarcraft" {`<br>`from corporate;`<br>`source any;`<br>`source-region any;`<br>`to internet;`<br>`destination any;`<br>`destination-region any;`<br>`user acme\duane;`<br>`category any;`<br>`application/service worldofwarcraft;`<br>`action deny;`<br>`terminal no;`<br>`}` |

 © Palo Alto Networks, Inc.

**Verify the User-ID Configuration (Continued)**

Step 4 Test your Captive Portal configuration.

1. From the same zone, go to a machine that is not a member of your directory, such as a Mac OS system, and try to ping to a system external to the zone. The ping should work without requiring authentication.
2. From the same machine, open a browser and navigate to a web site in a destination zone that matches a Captive Portal policy you defined. You should see the Captive Portal web form.

3. Log in using the correct credentials and confirm that you are redirected to the requested page.
4. You can also test your Captive Portal policy using the `test cp-policy-match` command as follows:

```
test cp-policy-match from corporate to internet
source 192.168.201.10 destination 8.8.8.8


Matched rule: 'captive portal' action: web-form
```

Step 5 Verify that user names are displayed in the log files (**Monitor > Logs**).

**Verify the User-ID Configuration (Continued)**

Step 6 Verify that user names are displayed in reports (**Monitor > Reports**). For example, when drilling down into the denied applications report, you should see a list of the users who attempted to access the applications as in the following example.

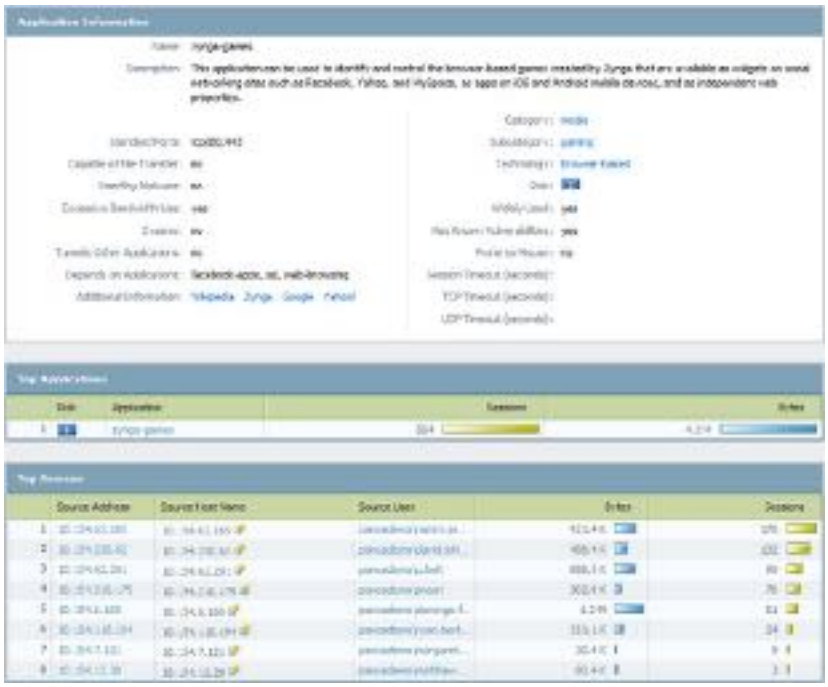

 © Palo Alto Networks, Inc.

# App-ID

To safely enable applications on your network, the Palo Alto Networks next-generation firewalls provide both an application and web perspective—App-ID and URL Filtering—to protect against a full spectrum of legal, regulatory, productivity, and resource utilization risks.

App-ID enables visibility into the applications on the network, so you can learn how they work, understand their behavioral characteristics and their relative risk. This application knowledge allows you to create and enforce security policies to enable, inspect, and shape desired applications and block undesired applications. When you define policies to begin allowing traffic, App-ID begins to classify traffic without any additional configuration.


- App-ID Overview
- Manage Custom or Unknown Applications
- Best Practices for Using App-ID in Policy
- Applications with Implicit Support
- About Application Level Gateways
- Disable the SIP Application-level Gateway (ALG)

# App-ID Overview

App-ID, a patented traffic classification system only available in Palo Alto Networks firewalls, determines what the application is irrespective of port, protocol, encryption (SSH or SSL) or any other evasive tactic used by the application. It applies multiple classification mechanisms—application signatures, application protocol decoding, and heuristics—to your network traffic stream to accurately identify applications.

Here's how App-ID identifies applications traversing your network:

- Traffic is matched against policy to check whether it is allowed on the network.
- Signatures are then applied to allowed traffic to identify the application based on unique application properties and related transaction characteristics. The signature also determines if the application is being used on its default port or it is using a non-standard port. If the traffic is allowed by policy, the traffic is then scanned for threats and further analyzed for identifying the application more granularly.
- If App-ID determines that encryption (SSL or SSH) is in use, and a decryption policy is in place, the session is decrypted and application signatures are applied again on the decrypted flow.
- Decoders for known protocols are then used to apply additional context-based signatures to detect other applications that may be tunneling inside of the protocol (e.g., Yahoo! Instant Messenger used across HTTP). Decoders validate that the traffic conforms to the protocol specification and provide support for NAT traversal and opening dynamic pinholes for applications such as SIP and FTP.
- For applications that are particularly evasive and cannot be identified through advanced signature and protocol analysis, heuristics or behavioral analysis may be used to determine the identity of the application.

When the application is identified, the policy check determines how to treat the application, for example—block, or allow and scan for threats, inspect for unauthorized file transfer and data patterns, or shape using QoS.

 © Palo Alto Networks, Inc.

# Manage Custom or Unknown Applications

Palo Alto Networks provides weekly App-ID updates to identify new applications. By default, App-ID is always enabled on the firewall, and you don't need to enable a series of signatures to identify well-known applications. Typically, the only applications that are classified as *unknown* traffic—*tcp, udp or non-syn-tcp*—in the ACC and the traffic logs are commercially available applications that have not yet been added to App-ID, internal or custom applications on your network, or potential threats.

On occasion, the firewall may report an application as unknown for the following reasons:

- Incomplete data—A handshake took place, but no data packets were sent prior to the timeout.
- Insufficient data—A handshake took place followed by one or more data packets; however, not enough data packets were exchanged to identify the application.

The following choices are available to handle unknown applications:

- Create security policies to control unknown applications by unknown TCP, unknown UDP or by a combination of source zone, destination zone, and IP addresses.
- Request an App-ID from Palo Alto Networks—If you would like to inspect and control the applications that traverse your network, for any unknown traffic, you can record a packet capture. If the packet capture reveals that the application is a commercial application, you can submit this packet capture to Palo Alto Networks for App-ID development. If it is an internal application, you can create a custom App-ID and/or define an application override policy.
- Create a custom App-ID with a signature and attach it to a security policy, or create a custom App-ID and define an application override policy—A custom App-ID allows you to customize the definition of the internal application—its characteristics, category and sub-category, risk, port, timeout—and exercise granular policy control in order to minimize the range of unidentified traffic on your network. Creating a custom App-ID also allows you to correctly identify the application in the ACC and traffic logs and is useful in auditing/reporting on the applications on your network. For a custom application you can specify a signature and a pattern that uniquely identifies the application and attach it to a security policy that allows or denies the application.

In order to collect the right data to create a custom application signature, you'll need a good understanding of packet captures and how datagrams are formed. If the signature is created too broadly you might inadvertently include other similar traffic; if it is defined too narrowly, the traffic will evade detection if it does not strictly match the pattern.

Custom App-IDs are stored in a separate database on the firewall and this database is not impacted by the weekly App-ID updates.

The supported application protocol decoders that enables the firewall to detect applications that may be tunneling inside of the protocol include the following as of content update 424: HTTP, HTTPS, DNS, FTP, IMAP SMTP, Telnet, IRC (Internet Relay Chat), Oracle, RTMP, RTSP, SSH, GNU-Debugger, GIOP (Global Inter-ORB Protocol), Microsoft RPC, Microsoft SMB (also known as CIFS). Additionally, with the 4.0 release of PAN-OS this custom App-ID capability has been expanded to include unknown TCP and unknown UDP as well.

Alternatively, if you would like the firewall to process the custom application using fast path (Layer-4 inspection instead of using App-ID for Layer-7 inspection), you can reference the custom App-ID in an application override policy. An application override with a custom application will prevent the session from being processed by the App-ID engine, which is a Layer-7 inspection. Instead it forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4, and thereby saves application processing time.

© Palo Alto Networks, Inc.

For example, if you build a custom application that triggers on a host header *www.mywebsite.com*, the packets are first identified as *web-browsing* and then are matched as your custom application (whose parent application is web-browsing). Because the parent application is web-browsing, the custom application is inspected at Layer-7 and scanned for content and vulnerabilities.

If you define an application override, the firewall stops processing at Layer-4. The custom application name is assigned to the session to help identify it in the logs, and the traffic is not scanned for threats.

For more details, refer to the following articles:

- Identifying Unknown Applications
- Video: How to Configure a Custom App-ID
- Custom Application Signatures

 © Palo Alto Networks, Inc.

# Best Practices for Using App-ID in Policy

- Review the ACC for the list of applications on your network, and determine which applications to allow or block. If you are migrating from a firewall where you defined port-based rules, to get a list of applications that run on a given port search by the port number in the application browser (**Objects > Applications**) on the Palo Alto Networks firewall or in Applipedia.
- Use **application-default** for the **Service**. The firewall compares the port used with the list of default ports for that application. If the port used is not a default port for the application, the firewall drops the session and logs the message **appid policy lookup deny**. If you have a application that is accessed on many ports and you would like to limit the ports on which the application is used, specify it in **Service**/**Service Group** objects in policies.
- Use application filters to dynamically include new applications in existing policy rules. See an example.

# Applications with Implicit Support

When creating a policy to allow specific applications, you must also be sure that you are allowing any other applications on which the application depends. In many cases, you do not have to explicitly allow access to the dependent applications in order for the traffic to flow because the firewall is able to determine the dependencies and allow them implicitly. This implicit support also applies to custom applications that are based on HTTP, SSL, MS-RPC, or RTSP. Applications for which the firewall cannot determine dependent applications on time will require that you explicitly allow the dependent applications when defining your policies. You can determine application dependencies in Applipedia.

The following table lists the applications for which the firewall has implicit support (as of Content Update 469).

**Table: Applications with Implicit Support**

| Application | Implicitly Supports |
|---|---|
| 360-safeguard-update | http |
| apple-update | http |
| apt-get | http |
| as2 | http |
| avg-update | http |
| avira-antivir-update | http, ssl |
| blokus | rtmp |
| bugzilla | http |
| clubcooee | http |
| corba | http |
| cubby | http, ssl |
| dropbox | ssl |
| esignal | http |
| evernote | http, ssl |
| ezhelp | http |
| facebook | http, ssl |
| facebook-chat | jabber |
| facebook-social-plugin | http |
| fastviewer | http, ssl |
| forticlient-update | http |
| gmail | http |
| good-for-enterprise | http, ssl |
| google-cloud-print | http, ssl, jabber |

 © Palo Alto Networks, Inc.

| Application | Implicitly Supports |
|---|---|
| google-desktop | http |
| google-drive-web | http |
| google-talk | jabber |
| google-update | http |
| gotomypc-desktop-sharing | citrix-jedi |
| gotomypc-file-transfer | citrix-jedi |
| gotomypc-printing | citrix-jedi |
| hipchat | http |
| iheartradio | ssl, http, rtmp |
| infront | http |
| instagram | http, ssl |
| issuu | http, ssl |
| java-update | http |
| jepptech-updates | http |
| kerberos | rpc |
| kik | http, ssl |
| lastpass | http, ssl |
| logmein | http, ssl |
| mcafee-update | http |
| megaupload | http |
| metatrader | http |
| mocha-rdp | t_120 |
| mount | rpc |
| ms-frs | msrpc |
| ms-rdp | t_120 |
| ms-scheduler | msrpc |
| ms-service-controller | msrpc |
| nfs | rpc |
| oovoo | http, ssl |
| paloalto-updates | ssl |
| panos-global-protect | http |
| panos-web-interface | http |

© Palo Alto Networks, Inc.

| Application | Implicitly Supports |
|---|---|
| pastebin | http |
| pastebin-posting | http |
| pinterest | http, ssl |
| portmapper | rpc |
| prezi | http, ssl |
| rdp2tcp | t_120 |
| renren-im | jabber |
| roboform | http, ssl |
| salesforce | http |
| stumbleupon | http |
| supremo | http |
| symantec-av-update | http |
| trendmicro | http |
| trillian | http, ssl |
| twitter | http |
| whatsapp | http, ssl |
| xm-radio | rtsp |


© Palo Alto Networks, Inc.

# About Application Level Gateways

The Palo Alto Networks firewall does not classify traffic by port and protocol; instead it identifies the application based on its unique properties and transaction characteristics using the App-ID technology. Some applications, however, require the firewall to dynamically open *pinholes* to establish the connection, determine the parameters for the session and negotiate the ports that will be used for the transfer of data; these applications use the application-layer payload to communicate the dynamic TCP or UDP ports on which the application opens data connections. For such applications, the firewall serves as an Application Level Gateway (ALG), and it opens a pinhole for a limited time and for exclusively transferring data or control traffic. The firewall also performs a NAT rewrite of the payload when necessary.

The Palo Alto Networks firewall functions as an ALG for the following protocols: FTP, SIP, H.323, RTSP, Oracle/SQLNet/TNS, MGCP, Unistim, SCCP, protocols.

When the firewall serves as an ALG for the Session Initiation Protocol (SIP), by default it performs NAT on the payload and opens dynamic pinholes for media ports. In some cases, depending on the SIP applications in use in your environment, the SIP endpoints have NAT intelligence embedded in their clients. In such cases, you might need to disable the SIP ALG functionality to prevent the firewall from modifying the signaling sessions. When SIP ALG is disabled, if App-ID determines that a session is SIP, the payload is not translated and dynamic pinholes are not opened. See Disable the SIP Application-level Gateway (ALG).

# Disable the SIP Application-level Gateway (ALG)

The Palo Alto Networks firewall uses the Session Initiation Protocol (SIP) application-level gateway (ALG) to open dynamic pinholes in the firewall where NAT is enabled. However, some applications—such as VoIP—have NAT intelligence embedded in the client application. In these cases, the SIP ALG on the firewall can interfere with the signaling sessions and cause the client application to stop working.

One solution to this problem is to define an Application Override Policy for SIP, but using this approach disables the App-ID and threat detection functionality. A better approach is to disable the SIP ALG, which does not disable App-ID or threat detection.

The following procedure describes how to disable the SIP ALG.

**Disable the SIP ALG**

Step 1 Select **Objects > Applications**.

Step 2 Select the **sip** application.

You can type **sip** in the **Search** box to help find the sip application.

Step 3 Select **Customize...** for **ALG** in the Options section of the Application dialog box.

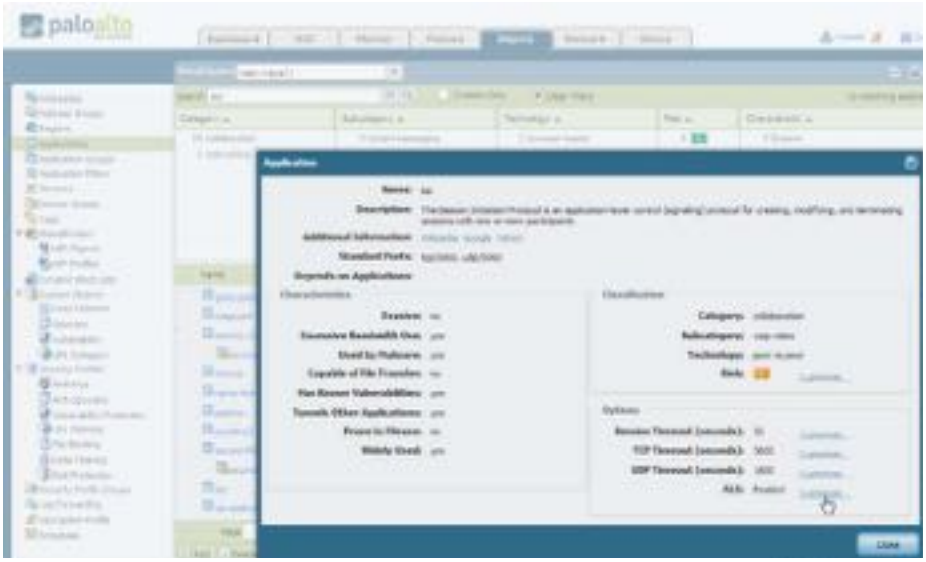

Step 4 Select the **Disable ALG** check box in the Application - sip dialog box and click **OK**.

Step 5 **Close** the Application dialog box and **Commit** the change.

 © Palo Alto Networks, Inc.

# Threat Prevention

The Palo Alto Networks next-generation firewall protects and defends your network from commodity threats and advanced persistent threats (APTs). The firewall's multi-pronged detection mechanisms include a signature-based (IPS/Command and Control/Antivirus) approach, heuristics-based (bot detection) approach, sandbox-based (WildFire) approach, and Layer 7 protocol analysis-based (App-ID) approach.

Commodity threats are exploits that are less sophisticated and more easily detected and prevented using a combination of the antivirus, anti-spyware, vulnerability protection and the URL filtering/Application identification capabilities on the firewall.

Advanced threats are perpetuated by organized cyber criminals or malicious groups that use sophisticated attack vectors to target your network, most commonly for intellectual property theft and financial data theft. These threats are more evasive and require intelligent monitoring mechanisms for detailed host and network forensics on malware. The Palo Alto Networks next-generation firewall in conjunction with WildFire and Panorama provides a comprehensive solution that intercepts and break the attack chain and provides visibility to prevent security infringement on your network—including mobile and virtualized—infrastructure.


- License the Threat Prevention Features
- About Security Profiles
- Set Up Security Profiles and Policies
- Prevent Brute Force Attacks
- Customize the Action and Trigger Conditions for a Brute Force Signature
- Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions
- Passive DNS Collection
- Content Delivery Network Infrastructure for Dynamic Updates
- Threat Prevention Resources


© Palo Alto Networks, Inc.

# License the Threat Prevention Features

The following sections describe the available licenses required to utilize threat prevention features and describes the activation process:

▲ About Threat Prevention Licenses

▲ Obtain and Install Licenses

## About Threat Prevention Licenses

The following is the list of licenses required for enabling all the threat prevention features on the firewall:

- **Threat Prevention**—Provides antivirus, anti-spyware, and vulnerability protection.
- **URL Filtering**—Provides the ability to control access to websites based on URL category. You can purchase and install a subscription for PAN-DB (Palo Alto Networks database) or the BrightCloud URL filtering databases.
- **WildFire**—The WildFire feature is included as part of the base product. This means that anyone can configure a file blocking profile to forward Portable Executable (PE) files to WildFire for analysis. A threat prevention subscription is required in order to receive antivirus signature updates, which include signatures discovered by WildFire.

  The WildFire subscription service provides enhanced services for organizations that require immediate security, enabling sub-hourly WildFire signature updates, advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet), as well as the ability to upload files using the WildFire API. A WildFire subscription is also required if your firewalls will be forwarding files to a private WF-500 WildFire appliance.
- **Decryption Port Mirror—**Provides the ability to create a copy of decrypted traffic from the firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as NetWitness or Solera—for archiving and analysis.Currently this license is available free of charge through the Palo Alto Networks Support Portal; this feature is supported on the PA-7050, PA-5000 Series and PA-3000 Series platforms only. For more information, see Configure Decryption Port Mirroring.

## Obtain and Install Licenses

To purchase licenses, contact the Palo Alto Networks sales department. After you obtain a license, navigate to **Device > Licenses**.

You can perform the following tasks depending on how you receive your licenses:

- **Retrieve license keys from license server**—Use this option if your license has been activated on the support portal.
- **Activate feature using authorization code**—Use this option to enable purchased subscriptions using an authorization code for licenses that have not been previously activated on the support portal.

 © Palo Alto Networks, Inc.

- **Manually upload license key**—Use this option if your device does not have connectivity to the Palo Alto Networks support site. In this case, you must download a license key file from the support site on an Internet connected computer and then upload to the device.

For more details about registering and activating licenses on your firewall, see Activate Licenses.

# About Security Profiles

Security profiles provide threat protection in security policies. For example, you can apply an antivirus profile to a security policy and all traffic that matches the security policy will be scanned for viruses.

For basic configuration examples to help you get started with these features, see Set Up Security Profiles and Policies.

The following table provides an overview of the security profiles that can be applied to security policies as well as a basic description of DoS and Zone Protection profiles:

| Threat Prevention Feature | Description |
|---|---|
| Antivirus | Protects against viruses, worms, and trojans as well as spyware downloads. Using a stream-based malware prevention engine, which inspects traffic the moment the first packet is received, the Palo Alto Networks antivirus solution can provide protection for clients without significantly impacting the performance of the firewall. This feature will scan for a wide variety of malware in executables, PDF files, HTML and JavaScript viruses, including support for scanning inside compressed files and data encoding schemes. Scanning of decrypted content can be performed by enabling decryption on the firewall.<br>The default profile inspects all of the listed protocol decoders for viruses, and generates alerts for SMTP, IMAP, and POP3 protocols while blocking for FTP, HTTP, and SMB protocols. Customized profiles can be used to minimize antivirus inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms.<br>The Palo Alto Networks WildFire system also provides signatures for persistent threats that are more evasive and have not yet been discovered by other antivirus solutions. As threats are discovered by WildFire, signatures are quickly created and then integrated into the standard antivirus signatures that can be downloaded by Threat Prevention subscribers on a daily basis (sub-hourly for WildFire subscribers). |

 © Palo Alto Networks, Inc.

| Threat Prevention Feature | Description |
|---|---|
| Anti-spyware | Blocks attempts from spyware on compromised hosts trying to phone-home or beacon out to external command-and-control (C2) servers. You can apply various levels of protection between zones. For example, you may want to have custom anti-spyware profiles that minimize inspection between trusted zones, while maximizing inspection on traffic received from an untrusted zone, such as Internet facing zones.<br>You can choose between two pre-defined profiles when applying anti-spyware to a security policy.<br>• **Default**—The default profile will use the default action for every signature, as specified by Palo Alto Networks when the signature is created.<br>• **Strict**—The strict profile will override the action of critical, high, and medium severity threats to the block action, regardless of the action defined in the signature file. The default action is taken with medium and informational severity signatures.<br>In PAN-OS 6.0, the DNS sinkhole feature has been added. The DNS sinkhole action that can be enabled in Anti-Spyware profiles enables the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. This feature can be used to identify infected hosts on the protected network using DNS traffic in situations where the firewall cannot see the infected client's DNS query (That is, the firewall cannot see the originator of the DNS query). In a typical deployment where the firewall is north of the local DNS server, the threat log will identify the local DNS resolver as the source of the traffic rather than the actual infected host. Sinkholing malware DNS queries solves this visibility problem by forging responses to the client host queries directed at malicious domains, so that clients attempting to connect to malicious domains (for command-and-control, for example) will instead attempt to connect to a sinkhole IP address you define. Infected hosts can then be easily identified in the traffic and threat logs because any host that attempts to connect to the sinkhole IP address are most likely infected with malware. Anti-spyware and vulnerability protection profiles are configured similarly. With anti-spyware, the main purpose is to detect malicious traffic leaving the network from infected clients, whereas vulnerability protection protects against threats entering the network. |
| Vulnerability Protection | Stops attempts to exploit system flaws or gain unauthorized access to systems. For example, this feature will protect against buffer overflows, illegal code execution, and other attempts to exploit system vulnerabilities. The default vulnerability protection profile protects clients and servers from all known critical, high, and medium-severity threats. You can also create exceptions, which allow you to change the response to a specific signature.<br>Anti-spyware and vulnerability protection profiles are configured similarly. With vulnerability protection, the main purpose is to detect malicious traffic entering the network from infected clients, whereas anti-spyware protection protects against threats leaving the network. |
| URL Filtering | Provides the ability to control user web traffic based on specific websites and/or website categories such as adult, shopping, gambling, and so on. The Palo Alto Networks PAN-DB URL database, or the BrightCloud database are available for categorization and enforcing URL Filtering policies on the firewall. |

| Threat Prevention Feature | Description |
|---|---|
| Data Filtering | Helps to prevent sensitive information such as credit card or social security numbers from leaving a protected network. You can also filter on key words, such as a sensitive project name and the word confidential. It is important to focus your profile on the desired file types to reduce false positives. For example, you may only want to search Word documents or Excel spreadsheets. You may also only want to scan web-browsing traffic, or FTP.<br><br>You can use default profiles, or create custom data patterns. There are two default profiles:<br>• **CC# (Credit Card)**—Identifies credit card numbers using a hash algorithm. The content must match the hash algorithm in order for data to be detected as a credit card number. This method will reduce false positives.<br>• **SSN# (Social Security Number)**—Uses an algorithm to detect nine digit numbers, regardless of format. There are two fields: **SSN#** and **SSN# (no dash)**.<br><br>Weight and Threshold Values<br><br>It is important to understand how the weight of an object (SSN, CC#, pattern) is calculated in order to set the appropriate threshold for a condition you are trying to filter. Each occurrence multiplied by the weight value will be added together in order to reach an action threshold (alert or block).<br><br>Example 1<br><br>For simplicity, if you only want to filter files with Social Security Numbers (SSN) and you define a weight of 3 for SSN#, you would use the following formula: each instance of a SSN x weight = threshold increment. In this case, if a Word document has 10 social security numbers you multiply that by the weight of 3, so 10 x 3 = 30. In order to take action for a file that contains 10 social security numbers you would set the threshold to 30. You may want to set an alert at 30 and then block at 60. You may also want to set a weight in the field SSN# (no dash) for Social Security Numbers that do not contain dashes. If multiple settings are used, they will accumulate to reach a given threshold.<br><br>Example 2<br><br>In this example, we will filter on files that contain Social Security Numbers and the custom pattern confidential. In other words, if a file has Social Security Numbers in addition to the word confidential and the combined instances of those items hit the threshold, the file will trigger an alert or block, depending on the action setting.<br>`SSN# weight = 3`<br>`Custom Pattern confidential weight = 20`<br> The custom pattern is case sensitive.<br><br>If the file contains 20 Social Security Numbers and a weight of 3 is configured, that is 20 x 3 = 60. If the file also contains one instance of the term confidential and a weight of 20 is configured, that is 1 x 20 = 20 for a total of 80. If your threshold for block is set to 80, this scenario would block the file. The alert or block action will be triggered as soon as the threshold is hit. |

 © Palo Alto Networks, Inc.

| Threat Prevention Feature | Description |
|---|---|
| File Blocking | Blocks specified file types over specified applications and in the specified session flow direction (inbound/outbound/both). You can set the profile to alert or block on upload and/or download and you can specify which applications will be subject to the file blocking profile. You can also configure custom block pages that will appear when a user attempts to download the specified file type. This allows the user to take a moment to consider whether or not they want to download a file.<br>The following actions can be set when the specified file is detected:<br>• **Alert**—When the specified file type is detected, a log is generated in the data filtering log.<br>• **Block**—When the specified file type is detected, the file is blocked and a customizable block page is presented to the user. A log is also generated in the data filtering log.<br>• **Continue**—When the specified file type is detected, a customizable continuation page is presented to the user. The user can click through the page to download the file. A log is also generated in the data filtering log.<br>• **Forward**—When the specified file type is detected, the file is sent to WildFire for analysis. A log is also generated in the data filtering log.<br>• **Continue-and-forward**—When the specified file type is detected, a customizable continuation page is presented to the user. The user can click through the page to download the file. If the user clicks through the continue page to download the file, the file is sent to WildFire for analysis. A log is also generated in the data filtering log. |
| DoS Protection | Provide detailed control for Denial of Service (DoS) protection policies. DoS policies allow you to control the number of sessions between interfaces, zones, addresses, and countries based on aggregate sessions or source and/or destination IP addresses. There are two DoS protection mechanisms that the Palo Alto Networks firewalls support.<br>• **Flood Protection**—Detects and prevents attacks where the network is flooded with packets resulting in too many half-open sessions and/or services being unable to respond to each request. In this case the source address of the attack is usually spoofed.<br>• **Resource Protection**— Detects and prevent session exhaustion attacks. In this type of attack, a large number of hosts (bots) are used to establish as many fully established sessions as possible to consume all of a system's resources.<br>Both of these protection mechanisms can be defined in a single DoS profile.<br>The DoS profile is used to specify the type of action to take and details on matching criteria for the DoS policy. The DoS profile defines settings for SYN, UDP, and ICMP floods, can enable resource protect and defines the maximum number of concurrent connections. After you configure the DoS protection profile, you then attach it to a DoS policy.<br>When configuring DoS protection, it is important to analyze your environment in order to set the correct thresholds and due to some of the complexities of defining DoS protection policies, this guide will not go into detailed examples. For more information, refer to the Threat Prevention Tech Note. |

| Threat Prevention Feature | Description |
|---|---|
| Zone Protection | Provides additional protection between specific network zones in order to protect the zones against attack. The profile must be applied to the entire zone, so it is important to carefully test the profiles in order to prevent issues that may arise with the normal traffic traversing the zones. When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session. For more information, refer to the Threat Prevention Tech Note. |

 © Palo Alto Networks, Inc.

# Set Up Security Profiles and Policies

The following sections provide basic threat prevention configuration examples:

- Set Up Antivirus, Anti-spyware, and Vulnerability Protection
- Set Up Data Filtering
- Set Up File Blocking

For information on controlling web access as part of your threat prevention strategy, see URL Filtering.

## Set Up Antivirus, Anti-spyware, and Vulnerability Protection

The following describes the steps needed to set up the default Antivirus, Anti-spyware, and Vulnerability Protection profiles. These features are very similar in purpose, so the following steps are just general steps needed to enable the default profiles.

All Anti-spyware and Vulnerability Protection signatures have a default action defined by Palo Alto Networks. You can view the default action by navigating to **Objects > Security Profiles > Anti-Spyware** or **Objects > Security Profiles >Vulnerability Protection** and then selecting a profile. Click the **Exceptions** tab and then click **Show all signatures** and you will see a list of the signatures with the default action in the Action column. To change the default action, you must create a new profile and then create rules with a non-default action, and/or add individual signature exceptions to **Exceptions** in the profile.

| **Set up Antivirus/Anti-spyware/Vulnerability Protection** | | |
|---|---|---|
| Step 1 | Verify that you have a Threat Prevention license. | • The Threat Prevention license bundles the Antivirus, Anti-spyware, and the Vulnerability Protection features in one license.<br>• Select **Device > Licenses** to verify that the **Threat Prevention** license is installed and check the expiration date. |
| Step 2 | Download the latest antivirus threat signatures. | 1. Select **Device > Dynamic Updates** and click **Check Now** at the bottom of the page to retrieve the latest signatures.<br>In the **Actions** column, click **Download** to install the latest Antivirus and Applications and Threats signatures. |

| Set up Antivirus/Anti-spyware/Vulnerability Protection (Continued) | |
|---|---|
| Step 3 Schedule signature updates. | 1. From **Device > Dynamic Updates**, click the text to the right of **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.<br>2. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select **Download Only**, you would need to manually go in and click the **Install** link in the **Action** column to install the signature. When you click **OK**, the update is scheduled. No commit is required.<br>3. (Optional) You can also enter the number of hours in the **Threshold** field to indicate the minimum age of a signature before a download will occur. For example, if you entered **10**, the signature must be at least 10 hours old before it will be downloaded, regardless of the schedule.<br>4. In an HA configuration, you can also click the **Sync To Peer** option to synchronize the content update with the HA peer after download/install. This will not push the schedule settings to the peer device, you need to configure the schedule on each device. |

### Best Practices for Antivirus Schedules

The general recommendation for antivirus signature update schedules is to perform a **download-and-install** on a daily basis for antivirus and weekly for applications and vulnerabilities.

#### Recommendations for HA Configurations:

- **Active/Passive HA**—If the MGT port is used for antivirus signature downloads, you should configure a schedule on both devices and both devices will download/install independently. If you are using a data port for downloads, the passive device will not perform downloads while it is in the passive state. In this case you would set a schedule on both devices and then select the **Sync To Peer** option. This will ensure that whichever device is active, the updates will occur and will then push to the passive device.
- **Active/Active HA**—If the MGT port is used for antivirus signature downloads on both devices, then schedule the download/install on both devices, but do not select the **Sync To Peer** option. If you are using a data port, schedule the signature downloads on both devices and select **Sync To Peer**. This will ensure that if one device in the active/active configuration goes into the active-secondary state, the active device will download/install the signature and will then push it to the active-secondary device.

 © Palo Alto Networks, Inc.

| Set up Antivirus/Anti-spyware/Vulnerability Protection (Continued) | |
|---|---|
| Step 4 Attach the security profiles to a security policy. | 1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab.<br>2. In **Profile Settings**, click the drop-down next to each security profile you would like to enable. In this example we choose default for **Antivirus, Vulnerability Protection,** and **Anti-spyware**.<br>If no security profiles have been previously defined, select **Profiles** from the **Profile Type** drop-down. You will then see the list of options to select the security profiles.<br> 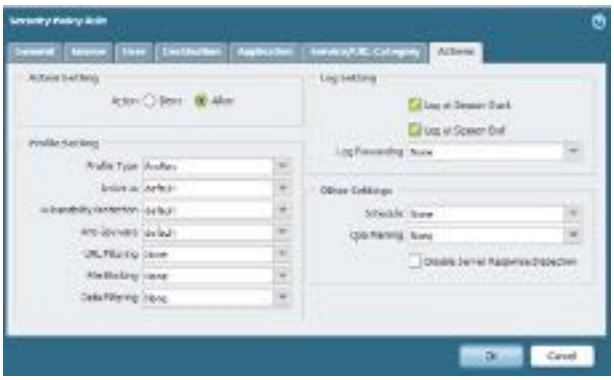 |
| Step 5 Save the configuration. | Click **Commit**. |

## Set Up Data Filtering

The following describes the steps needed to configure a data filtering profile that will detect Social Security Numbers and a custom pattern identified in .doc and .docx documents.

| Data Filtering Configuration Example | | |
|---|---|---|
| Step 1 | Create a Data Filtering security profile. | 1. Select **Objects > Security Profiles > Data Filtering** and click **Add**.<br>2. Enter a **Name** and a **Description** for the profile. In this example the name is *DF_Profile1* with the description *Detect Social Security Numbers*.<br>3. (Optional) If you want to collect data that is blocked by the filter, select the **Data Capture** check box.<br> You must set a password as described in Step 2 if you are using the data capture feature. |
| Step 2 | (Optional) Secure access to the data filtering logs to prevent other administrators from viewing sensitive data.<br>When you enable this option, you will be prompted for the password when you view logs in **Monitor > Logs > Data Filtering**. | 1. Select **Device > Setup > Content-ID**.<br>2. Click **Manage Data Protection** in the Content-ID Features section.<br>3. Set the password that will be required to view the data filtering logs. |

 © Palo Alto Networks, Inc.

**Data Filtering Configuration Example (Continued)**

Step 3 Define the data pattern that will be used in the Data Filtering Profile.

In this example, we will use the keyword **confidential** and will set the option to search for SSN numbers with dashes (Example - 987-654-4320).

It is helpful to set the appropriate thresholds and define keywords within documents to reduce false positives.

1. From the Data Filtering Profile page click **Add** and select **New** from the **Data Pattern** drop-down. You can also configure data patterns from **Objects > Custom Signatures > Data Patterns**.
2. For this example, name the Data Pattern signature **Detect SS Numbers** and add the description **Data Pattern to detect Social Security numbers**.
3. In the **Weight** section for **SSN#** enter **3**. See Weight and Threshold Values for more details.

4. (Optional) You can also set **Custom Patterns** that will be subject to this profile. In this case, you specify a pattern in the custom patterns **Regex** field and set a weight. You can add multiple match expressions to the same data pattern profile. In this example, we will create a **Custom Pattern** named **SSN_Custom** with a custom pattern of **confidential** (the pattern is case sensitive) and use a weight of **20**. The reason we use the term **confidential** in this example is because we know that our social security Word docs contain this term, so we define that specifically.

| Data Filtering Configuration Example (Continued) | | |
|---|---|---|
| Step 4 | Specify which applications to filter and set the file types. | 1. Set **Applications** to **Any.** This will detect any supported application such as: web-browsing, FTP, or SMTP. If you want to narrow down the application, you can select it from the list. For applications such as Microsoft Outlook Web App that uses SSL, you will need to enable decryption. Also make sure you understand the naming for each application. For example, Outlook Web App, which is the Microsoft name for this application is identified as the application **outlook-web** in the PAN-OS list of applications. You can check the logs for a given application to identify the name defined in PAN-OS.<br>2. Set **File Types** to **doc** and **docx** to only scan doc and docx files. |
| Step 5 | Specify the direction of traffic to filter and the threshold values. | 1. Set the **Direction** to **Both**. Files that are uploaded or downloaded will be scanned.<br>2. Set the **Alert Threshold** to **35.** In this case, an alert will be triggered if 5 instances of Social Security Numbers exist and 1 instance of the term **confidential** exists. The formula is 5 SSN instances with a weight of 3 = 15 plus 1 instance of the term confidential with a weight of 20 = 35.<br>3. Set the **Block Threshold** to **50**. The file will be blocked if the threshold of 50 instances of a SSN and/or the term confidential exists in the file. In this case, if the doc contained 1 instance of the word **confidential** with a weight of 20 that equals 20 toward the threshold, and the doc has 15 Social Security Numbers with a weight of 3 that equals 45. Add 20 and 45 and you have 65, which will exceed the block threshold of 50. |

 © Palo Alto Networks, Inc.

| Data Filtering Configuration Example (Continued) | | |
|---|---|---|
| Step 6 | Attach the Data Filtering profile to the security rule. | 1. Select **Policies > Security** and select the security policy rule to which to apply the profile.<br>2. Click the security policy rule to modify it and then click the **Actions** tab. In the **Data Filtering** drop-down, select the new data filtering profile you created and then click **OK** to save. In this example, the data filtering rule name is **DF_Profile1**.<br> 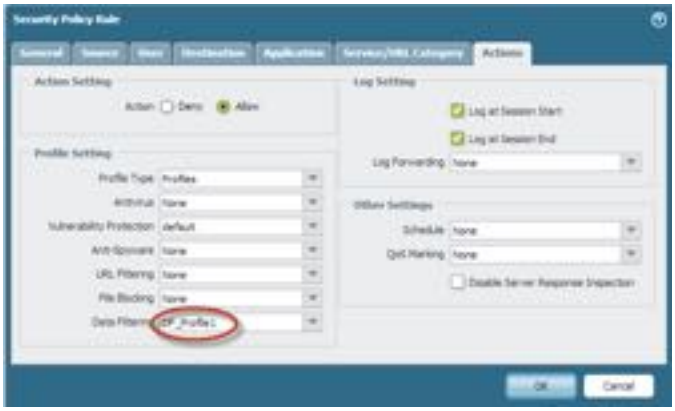 |
| Step 7 | **Commit** the configuration. | |
| Step 8 | Test the data filtering configuration.<br><br>If you have problems getting Data Filtering to work, you can check the Data Filtering log or the Traffic log to verify the application that you are testing with and make sure your test document has the appropriate number of unique Social Security Number instances. For example, an application such as Microsoft Outlook Web App may seem to be identified as web-browsing, but if you look at the logs, the application is **outlook-web**. Also increase the number of SSNs, or your custom pattern to make sure you are hitting the thresholds. | When testing, you must use real Social Security Numbers and each number must be unique. Also, when defining Custom Patterns as we did in this example with the word **confidential**, the pattern is case sensitive. To keep your test simple, you may want to just test using a data pattern first, then test the SSNs.<br>1. Access a client PC in the trust zone of the firewall and send an HTTP request to upload a .doc or .docx file that contains the exact information you defined for filtering.<br>2. Create a Microsoft Word document with one instance of the term confidential and five Social Security numbers with dashes.<br>3. Upload the file to a website. Use an HTTP site unless you have decryption configured, in which case you can use HTTPS.<br>4. Select **Monitoring > Logs > Data Filtering** logs.<br>5. Locate the log that corresponds to the file you just uploaded. To help filter the logs, use the source of your client PC and the destination of the web server. The action column in the log will show **reset-both**. You can now increase the number of Social Security Numbers in the document to test the block threshold.<br> |

# Set Up File Blocking

This example will describe the basic steps needed to set up file blocking and forwarding. In this configuration, we will configure the options needed to prompt users to continue before downloading .exe files from websites. When testing this example, be aware that you may have other systems between you and the source that may be blocking content.

| **Configure File Blocking** | | |
|---|---|---|
| Step 1 | Create the file blocking profile. | 1. Select **Objects > Security Profiles > File Blocking** and click **Add**.<br>2. Enter a **Name** for the file blocking profile, for example *Block_EXE*. Optionally enter a **Description**, such as *Block users from downloading exe files from websites.* |
| Step 2 | Configure the file blocking options. | 1. Click **Add** to define the profile settings.<br>2. Enter a **Name**, such as *BlockEXE*.<br>3. Set the **Applications** for filtering, for example **web-browsing**.<br>4. Set **File Types** to **exe**.<br>5. Set the **Direction** to **download**.<br>6. Set the **Action** to **continue**. By choosing the continue option, users will be prompted with a response page prompting them to click continue before the file will be downloaded.<br> <br>7. Click **OK** to save the profile. |
| Step 3 | Apply the file blocking profile to a security policy. | 1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Set Up Basic Security Policies.<br>2. Click the **Actions** tab within the policy rule.<br>3. In the Profile Settings section, click the drop-down and select the file blocking profile you configured. In this case, the profile name is *Block_EXE*.<br>4. **Commit** the configuration.<br>If no security profiles have been previously defined, select the **Profile Type** drop-down and select **Profiles**. You will then see the list of options to select the security profiles. |

 © Palo Alto Networks, Inc.

**Configure File Blocking (Continued)**

Step 4 To test your file blocking configuration, access a client PC in the trust zone of the firewall and attempt to download a .exe file from a website in the untrust zone. A response page should display. Click **Continue** to download the file. You can also set other actions, such as alert only, forward (which will forward to WildFire), or block, which will not provide a continue page to the user. The following shows the default response page for File Blocking:

**Example: Default File Blocking Response Page**

Step 5 (Optional) Define custom file blocking response pages (**Device > Response Pages**). This allows you to provide more information to users when they see a response page. You can include information such as company policy information and contact information for a Helpdesk.

When you create a file blocking profile with the action continue or continue-and-forward (used for WildFire forwarding), you can only choose the application web-browsing. If you choose any other application, traffic that matches the security policy will not flow through the firewall due to the fact that the users will not be prompted with a continue page. Also, if the website uses HTTPS, you will need to have a decryption policy in place.

You may want to check your logs to confirm what application is being used when testing this feature. For example, if you are using Microsoft Sharepoint to download files, even though you are using a web-browser to access the site, the application is actually `sharepoint-base`, or `sharepoint-document`. You may want to set the application type to **Any** for testing.

# Prevent Brute Force Attacks

A brute force attack uses a large volume of requests/responses from the same source or destination IP address to break into a system. The attacker employs a trial-and-error method to guess the response to a challenge or a request.

The Vulnerability Protection profile on the firewall includes signatures to protect you from brute force attacks. Each signature has an ID, Threat Name, Severity and is triggered when a pattern is recorded. The pattern specifies the conditions and interval at which the traffic is identified as a brute-force attack; some signatures are associated with another child signature that is of a lower severity and specifies the pattern to match against. When a pattern matches against the signature or child signature, it triggers the default action for the signature.

To enforce protection:

- Attach the vulnerability profile to a security rule. See Set Up Antivirus, Anti-spyware, and Vulnerability Protection.
- Install content updates that include new signatures to protect against emerging threats. See Manage Content Updates.

 © Palo Alto Networks, Inc.

# Customize the Action and Trigger Conditions for a Brute Force Signature

The firewall includes two types of predefined brute force signatures—parent signature and child signature. A child signature is a single occurrence of a traffic pattern that matches the signature. A parent signature is associated with a child signature and is triggered when multiple events occur within a time interval and match the traffic pattern defined in the child signature.

Typically, a child signature is of default action *allow* because a single event is not indicative of an attack. In most cases, the action for a child signature is set to allow so that legitimate traffic is not blocked and threat logs are not generated for non-noteworthy events. Therefore, Palo Alto Networks recommends that you only change the default action after careful consideration.

In most cases, the brute force signature is a noteworthy event because of its recurrent pattern. If you would like to customize the action for a brute-force signature, you can do one of the following:

- Create a rule to modify the default action for all signatures in the brute force category. You can define the action to allow, alert, block, reset, or drop the traffic.
- Define an exception for a specific signature. For example, you can search for a CVE and define an exception for it.

  For a parent signature, you can modify both the trigger conditions and the action; for a child signature only the action can be modified.

To effectively mitigate an attack, the **block-ip address** action is recommended over the drop or reset action for most brute force signatures.

| **Customize the Threshold and Action for a Signature** | | |
|---|---|---|
| Step 1 | Create a new Vulnerability Protection Profile. | 1. Select **Objects > Security Profiles > Vulnerability Protection**.<br>2. Click **Add** and enter a **Name** for the Vulnerability Protection Profile. |
| Step 2 | Create a rule that defines the action for all signatures in a category.<br> | 1. Select **Rules**, click **Add** and enter a **Name** for the rule.<br>2. Set the **Action**. In this example, it is set to **Block**.<br>3. Set **Category** to **brute-force**.<br>4. (Optional) If blocking, specify whether to block server or client, the default is any.<br>5. See Step 3 to customize the action for a specific signature.<br>6. See Step 4 to customize the trigger threshold for a parent signature.<br>7. Click **OK** to save the rule and the profile. |

| Customize the Threshold and Action for a Signature | |
|---|---|
| Step 3 (Optional) Customize the action for a specific signature.<br> <br> | 1. Select **Exceptions** and click **Show all signatures** to find the signature you want to modify.<br>To view all the signatures in the brute-force category, search for ( category contains 'brute-force' ).<br>2. To edit a specific signature, click the predefined default action in the **Action** column.<br>3. Set the action to **allow**, **alert** or **block-ip**.<br>4. If you select block-ip, complete these additional tasks:<br>a. Specify the **Time** period (in seconds) after which to trigger the action.<br>b. In the **Track By** field, define whether to block the IP address by **IP source** or by **IP source and destination**.<br> 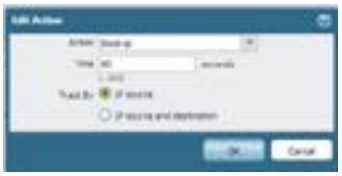 <br>5. Click **OK**<br>6. For each modified signature, select the checkbox in the **Enable** column.<br>7. Click **OK**. |
| Step 4 Customize the trigger conditions for a parent signature.<br>A parent signature that can be edited is marked with this icon .<br>In this example, the search criteria was brute force category and CVE-2008-1447. | 1. Click to edit the time attribute and the aggregation criteria for the signature.<br> <br>2. To modify the trigger threshold specify the **Number of Hits** *per x* **seconds**.<br>3. Specify whether to aggregate the number of hits by **source**, **destination** or by **source and destination**.<br>4. Click **OK**. |
| Step 5 Attach this new profile to a security rule. | |
| Step 6 Save your changes. | 1. Click **Commit**. |


© Palo Alto Networks, Inc.

# Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions

To monitor and protect your network from most Layer 4 and Layer 7 attacks, here are a few recommendations.

- Upgrade to the most current PAN-OS software version and content release version to ensure that you have the latest security updates. See Manage Content Updates and Install Software Updates.
- For web servers, create a security policy to only allow the protocols that the server supports. For example, ensure that only HTTP traffic is allowed to a web server. If you have defined an application override policy for a custom application, make sure to restrict access to specific source zone or set of IP addresses.
- Attach the following security profiles to your security policies to provide signature-based protection.
  - Create a vulnerability protection profile to block all vulnerabilities with severity low and higher.
  - Create a anti-spyware profile to block all spyware.
  - Create an antivirus profile to block all content that matches an antivirus signature.
- Block all unknown applications/traffic using security policy. Typically, the only applications that are classified as unknown traffic are internal or custom applications on your network, or potential threats. Because unknown traffic can be a non-compliant application or protocol that is anomalous or abnormal, or a known application that is using non-standard ports, unknown traffic should be blocked. See Manage Custom or Unknown Applications.
- Create a file blocking profile that blocks Portable Executable (PE) file types for Internet-based SMB (Server Message Block) traffic from traversing the trust to untrust zones, (ms-ds-smb applications).
- Create a zone protection profile that is configured to drop mismatched and overlapping TCP segments, to protect against packet-based attacks.

  By deliberately constructing connections with overlapping but different data in them, attackers can attempt to cause misinterpretation of the intent of the connection. This can be used to deliberately induce false positives or false negatives. An attacker can use IP spoofing and sequence number prediction to intercept a user's connection and inject his/her own data into the connection. PAN-OS uses this field to discard such frames with mismatched and overlapping data. The scenarios where the received segment will be discarded are:
  - The segment received is contained within another segment.
  - The segment received overlaps with part of another segment.
  - The segment completely contains another segment.
- Verify that support for IPv6 is enabled, if you have configured IPv6 addresses on your network hosts. (**Network > Interfaces > Ethernet> IPv6**)

  This allows access to IPv6 hosts and filters IPv6 packets that are encapsulated in IPv4 packets. Enabling support for IPv6 prevents IPv6 over IPv4 multicast addresses from being leveraged for network reconnaissance.

© Palo Alto Networks, Inc.

- Enable support for multicast traffic so that the firewall can enforce policy on multicast traffic. (**Network > Virtual Router > Multicast)**
- Enable the following CLI command to clear the URG bit flag in the TCP header and disallow out-of-band processing of packets.

  The urgent pointer in the TCP header is used to promote a packet for immediate processing by removing it from the processing queue and expediting it through the TCP/IP stack on the host. This process is called out-of-band processing. Because the implementation of the urgent pointer varies by host, to eliminate ambiguity, use the following CLI command to disallow out-of-band processing; the out-of-band byte in the payload becomes part of the payload and the packet is not processed urgently. Making this change allows you to remove ambiguity in how the packet is processed on the firewall and the host, and the firewall sees the exact same stream in the protocol stack as the host for whom the packet is destined.

  ```
  set deviceconfig setting tcp urgent-data clear
  ```

- Enable the following CLI command for disabling the bypass-exceed-queue.

  The bypass exceed queue is required for out of order packets. This scenario is most common in an asymmetric environment where the firewall receives packets out of order. For identification of certain applications (App-ID) the firewall performs heuristic analysis. If the packets are received out of order, the data must be copied to a queue in order to complete the analysis for the application.

  ```
  set deviceconfig setting application bypass-exceed-queue no
  ```

- Enable the following CLI commands for disabling the inspection of packets when the out-of-order packet limit is reached. The Palo Alto Networks firewall can collect up to 64 out-of-order packets per session. This counter identifies that packets have exceeded the 64-packet limit. When the bypass setting is set to **no**, the device drops the out-of-order packets that exceed the 64-packet limit. A commit is required.

  ```
  set deviceconfig setting tcp bypass-exceed-oo-queue no
  set deviceconfig setting ctd bypass-exceed-queue no
  ```

- Enable the following CLI commands for checking the TCP timestamp. The TCP timestamp records when the segment was sent and allows the firewall to verify that the timestamp is valid for that session.

  ```
  set deviceconfig setting tcp check-timestamp-option yes
  ```

- Disable the HTTP Range option. The HTTP Range option allows a client to fetch part of a file only. When a next-generation firewall in the path of a transfer identifies and drops a malicious file, it terminates the TCP session with a RST packet. If the web browser implements the HTTP Range option, it can start a new session to fetch only the remaining part of the file. This prevents the firewall from triggering the same signature again due to the lack of context into the initial session, while at the same time allowing the web browser to reassemble the file and deliver the malicious content. To prevent this, disable the HTTP Range option as follows:

  ```
  set deviceconfig setting ctd skip-block-http-range no
  ```

 © Palo Alto Networks, Inc.

# Passive DNS Collection

Passive DNS is an opt-in feature that enables the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis in order to improve threat intelligence and threat prevention capabilities. The data collected includes non-recursive (i.e. originating from the local recursive resolver, not individual clients) DNS query and response packet payloads. Data submitted via the Passive DNS Monitoring feature consists solely of mappings of domain names to IP addresses. Palo Alto Networks retains no record of the source of this data and does not have the ability to associate it with the submitter at a future date.

The Palo Alto Networks threat research team uses this information to gain insight into malware propagation and evasion techniques that abuse the DNS system. Information gathered through this data collection is used to improve accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control signatures, and WildFire.

DNS responses are only forwarded to the Palo Alto Networks and will only occur when the following requirements are met:

- DNS response bit is set
- DNS truncated bit is not set
- DNS recursive bit is not set
- DNS response code is 0 or 3 (NX)
- DNS question count bigger than 0
- DNS Answer RR count is bigger than 0 or if it is 0, the flags need to be 3 (NX)
- DNS query record type are A, NS, CNAME, AAAA, MX

Passive DNS monitoring is disabled by default, but it is recommended that you enable it to facilitate enhanced threat intelligence. Use the following procedure to enable Passive DNS:

**Enable Passive DNS**

1. Select **Objects > Security Profiles > Anti-Spyware**.
2. Select an existing profile to modify it or configure a new profile.

   The Anti-Spyware profile must be attached to a security policy that governs your DNS server's external DNS traffic.
3. Select the **DNS Signatures** tab and click the **Enable Passive DNS Monitoring** check box.
4. Click **OK** and then **Commit**.

# Content Delivery Network Infrastructure for Dynamic Updates

Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to the Palo Alto Networks devices. The devices access the web resources in the CDN to perform various App-ID and Content-ID functions. For enabling and scheduling the content updates, see Manage Content Updates.

The following table lists the web resources that the firewall accesses for a feature or application:

| Resource | URL | Static Addresses (If a static server is required) |
|---|---|---|
| Application Database | • updates.paloaltonetworks.com:443 | staticupdates.paloaltonetworks.com or the IP address 199.167.52.15 |
| Threat/Antivirus Database | • updates.paloaltonetworks.com:443<br>• downloads.paloaltonetworks.com:443<br>As a best practice, set the update server to updates.paloaltonetworks.com.This allows the Palo Alto Networks device to receive content updates from the server closest to it in the CDN infrastructure. | staticupdates.paloaltonetworks.com or the IP address 199.167.52.15 |
| PAN-DB URL Filtering | *.urlcloud.paloaltonetworks.com<br>Resolves to the primary URL s0000.urlcloud.paloaltonetworks.com and is then redirected to the regional server that is closest:<br>• s0100.urlcloud.paloaltonetworks.com<br>• s0200.urlcloud.paloaltonetworks.com<br>• s0300.urlcloud.paloaltonetworks.com<br>• s0500.urlcloud.paloaltonetworks.com | Static IP addresses are not available. However, you can manually resolve a URL to an IP address and allow access to the regional server IP address. |
| BrightCloud URL Filtering | • database.brightcloud.com:443/80<br>• service.brightcloud.com:80 | Contact BrightCloud Customer Support. |

 © Palo Alto Networks, Inc.

| Resource | URL | Static Addresses (If a static server is required) |
|---|---|---|
| WildFire | • beta.wildfire.paloaltonetworks.com:443/80<br>• beta-s1.wildfire.paloaltonetworks.com:443/80<br>**Note** Beta sites are only accessed by a firewall running a Beta release version.<br>• mail.wildfire.paloaltonetworks.com:25<br>• wildfire.paloaltonetworks.com:443/80 | • mail.wildfire.paloaltonetworks.com:25 or the IP address 54.241.16.83<br>• wildfire.paloaltonetworks.com:443/80 or 54.241.8.199<br>The regional URL/IP addresses are as follows:<br>• ca-s1.wildfire.paloaltonetworks.com:44 or 54.241.34.71<br>• va-s1.wildfire.paloaltonetworks.com:443 or 174.129.24.252<br>• eu-s1.wildfire.paloaltonetworks.com:443 or 54.246.95.247<br>• sg-s1.wildfire.paloaltonetworks.com:443 or 54.251.33.241<br>• jp-s1.wildfire.paloaltonetworks.com:443 or 54.238.53.161<br>• portal3.wildfire.paloaltonetworks.com:443/80 or 54.241.8.199<br>• ca-s3.wildfire.paloaltonetworks.com:443 or 54.241.34.71<br>• va-s3.wildfire.paloaltonetworks.com:443 or 23.21.208.35<br>• eu-s3.wildfire.paloaltonetworks.com:443 or 54.246.95.247<br>• sg-s3.wildfire.paloaltonetworks.com:443 or 54.251.33.241<br>• jp-s3.wildfire.paloaltonetworks.com:443 or 54.238.53.161<br>• wildfire.paloaltonetworks.com.jp:443/80 or 180.37.183.53<br>• wf1.wildfire.paloaltonetowrks.jp:443 or 180.37.180.37<br>• wf2.wildfire.paloaltonetworks.jp:443 or 180.37.181.18<br>• portal3.wildfire.paloaltonetworks.jp:443/80 or 180.37.183.53 |

© Palo Alto Networks, Inc.

# Threat Prevention Resources

For more information on Threat Prevention, refer to the following sources:

- Creating Custom Threat Signatures
- Threat Prevention Deployment
- Understanding DoS Protection

To view a list of Threats and Applications that Palo Alto Networks products can identify, use the following links:

- Applipedia—Provides details on the applications that Palo Alto Networks can identify.
- Threat Vault—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

 © Palo Alto Networks, Inc.

# Decryption

Palo Alto Networks firewalls provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Decryption on a Palo Alto Networks firewall includes the capability to enforce security policies on encrypted traffic, where otherwise the encrypted traffic might not be blocked and shaped according to your configured security settings. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted traffic. Enabling decryption on a Palo Alto Networks firewall can include preparing the keys and certificates required for decryption, creating a decryption policy, and configuring decryption port mirroring. See the following topics to learn about and configure decryption:

- Decryption Overview
- Decryption Concepts
- Configure SSL Forward Proxy
- Configure SSL Inbound Inspection
- Configure SSH Proxy
- Configure Decryption Exceptions
- Configure Decryption Port Mirroring

# Decryption Overview

Secure Sockets Layer (SSL) and Secure Shell (SSH) are encryption protocols used to secure traffic between two entities, such as a web server and a client. SSL and SSH encapsulate traffic, encrypting data so that it is meaningless to entities other than the client and server with the keys to decode the data and the certificates to affirm trust between the devices. Traffic that has been encrypted using the protocols SSL and SSH can be decrypted to ensure that these protocols are being used for the intended purposes only, and not to conceal unwanted activity or malicious content.

Palo Alto Networks firewalls decrypt encrypted traffic by using keys to transform strings (passwords and shared secrets) from ciphertext to plaintext (decryption) and from plaintext back to ciphertext (re-encrypting traffic as it exits the device). Certificates are used to establish the firewall as a trusted third party and to create a secure connection. SSL decryption (both forward proxy and inbound inspection) requires certificates to establish trust between two entities in order to secure an SSL/TLS connection. Certificates can also be used when excluding servers from SSL decryption. You can integrate a hardware security module (HSM) with a firewall to enable enhanced security for the private keys used in SSL forward proxy and SSL inbound inspection decryption. To learn more about storing and generating keys using an HSM and integrating an HSM with your firewall, see Secure Keys with a Hardware Security Module. SSH decryption does not require certificates.

Palo Alto Networks firewall decryption is policy-based, and can be used to decrypt, inspect, and control both inbound and outbound SSL and SSH connections. Decryption policies allow you to specify traffic for decryption according to destination, source, or URL category and in order to block or restrict the specified traffic according to your security settings. The firewall uses certificates and keys to decrypt the traffic specified by the policy to plaintext, and then enforces App-ID and security settings on the plaintext traffic, including Decryption, Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles. After traffic is decrypted and inspected on the firewall, the plaintext traffic is re-encrypted as it exits the firewall to ensure privacy and security. Use policy-based decryption on the firewall to achieve outcomes such as the following:

- Prevent malware concealed as encrypted traffic from being introduced into an corporate network.
- Prevent sensitive corporate information from moving outside the corporate network.
- Ensure the appropriate applications are running on a secure network.
- Selectively decrypt traffic; for example, exclude traffic for financial or healthcare sites from decryption by configuring a decryption exception.

The three decryption policies offered on the firewall, SSL Forward Proxy, SSL Inbound Inspection, and SSH Proxy, all provide methods to specifically target and inspect SSL outbound traffic, SSL inbound traffic, and SSH traffic, respectively. The decryption policies provide the settings for you to specify what traffic to decrypt and decryption profiles can be selected when creating a policy, in order to apply more granular security settings to decrypted traffic, such as checks for server certificates, unsupported modes, and failures. This policy-based decryption on the firewall gives you visibility into and control of SSL and SSH encrypted traffic according to configurable parameters.

You can also choose to extend a decryption configuration on the firewall to include Decryption Port Mirroring, which allows for decrypted traffic to be forwarded as plaintext to a third party solution for additional analysis and archiving.

 © Palo Alto Networks, Inc.

# Decryption Concepts

To learn about keys and certificates for decryption, decryption policies, and decryption port mirroring, see the following topics:

- Keys and Certificates for Decryption Policies
- SSL Forward Proxy
- SSL Inbound Inspection
- SSH Proxy
- Decryption Exceptions
- Decryption Port Mirroring

## Keys and Certificates for Decryption Policies

Keys are strings of numbers that are typically generated using a mathematical operation involving random numbers and large primes. Keys are used to transform other strings—such as passwords and shared secrets—from plaintext to ciphertext (called *encryption*) and from ciphertext to plaintext (called *decryption*). Keys can be symmetric (the same key is used to encrypt and decrypt) or asymmetric (one key is used for encryption and a mathematically related key is used for decryption). Any system can generate a key.

X.509 certificates are used to establish trust between a client and a server in order to establish an SSL connection. A client attempting to authenticate a server (or a server authenticating a client) knows the structure of the X.509 certificate and therefore knows how to extract identifying information about the server from fields within the certificate, such as its FQDN or IP address (called a *common name* or *CN* within the certificate) or the name of the organization, department, or user to which the certificate was issued. All certificates must be issued by a certificate authority (CA). After the CA verifies a client or server, the CA issues the certificate and signs it using its private key. With a decryption policy configured, an SSL/TLS session between the client and the server is only established if the firewall trusts the CA that signed the server's certificate. In order to establish trust, the firewall must have the server's root CA certificate in its certificate trust list (CTL) and use the public key contained in that root CA certificate to verify the signature. The firewall then presents a copy of the server certificate signed by the Forward Trust certificate for the client to authenticate. You can also configure the firewall to use an enterprise CA as a forward trust certificate for SSL Forward Proxy. If the firewall does not have the server's root CA certificate in it's CTL, the firewall will present a copy of the server certificate signed by the Forward Untrust certificate to the client. The Forward Untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites hosted by a server with untrusted certificates.

For detailed information on certificates, see Certificate Management.

To control the trusted CAs that your device trusts, use the **Device > Certificate Management > Certificates > Default Trusted Certificate Authorities** tab on the firewall's web interface.

Table: Palo Alto Networks Device Keys and Certificates describes the different keys and certificates used by Palo Alto Networks devices for decryption. As a best practice, use different keys and certificates for each usage.

**Table: Palo Alto Networks Device Keys and Certificates**

| Key/Certificate Usage | Description |
|---|---|
| Forward Trust | The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall trusts. To configure a Forward Trust certificate on the firewall, see Step 2 in the Configure SSL Forward Proxy task. For added security, the forward trust certificate can be stored on a Hardware Security Module (HSM), see Store Private Keys on an HSM. |
| Forward Untrust | The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall does not trust. To configure a Forward Untrust certificate on the firewall, see Step 3 in the Configure SSL Forward Proxy task. |

 © Palo Alto Networks, Inc.

| Key/Certificate Usage | Description |
|---|---|
| SSL Exclude Certificate | Certificates for servers that you want to exclude from SSL decryption. For example, if you have SSL decryption enabled, but have certain servers that you do not want included in SSL decryption, such as the web services for your HR systems, you would import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Exclude a Server From Decryption. |
| SSL Inbound Inspection | The certificate used to decrypt inbound SSL traffic for inspection and policy enforcement. For this application, you would import the server certificate for the servers for which you are performing SSL inbound inspection, or store them on an HSM (see Store Private Keys on an HSM). |

## SSL Forward Proxy

Use an SSL Forward Proxy decryption policy to decrypt and inspect SSL/TLS traffic from internal users to the web. SSL Forward Proxy decryption prevents malware concealed as SSL encrypted traffic from being introduced to your corporate network; for example, if an employee is using her Gmail account from her corporate office and opens an email attachment that contains a virus, SSL Forward Proxy decryption will prevent the virus from infecting the client system and entering the corporate network.

With SSL Forward Proxy decryption, the firewall resides between the internal client and outside server. The firewall uses Forward Trust or Forward Untrust certificates to establish itself a trusted third party to the session between the client and the server (For details on certificates, see Keys and Certificates for Decryption Policies). When the client initiates an SSL session with the server, the firewall intercepts the client's SSL request and forwards the SSL request to the server. The server sends a certificate intended for the client that is intercepted by the firewall. If the server's certificate is signed by a CA that the firewall trusts, the firewall creates a copy of the server's certificate signed by the Forward Trust certificate and sends the certificate to the client to authenticate. If the server's certificate is signed by a CA that the firewall does not trust, the firewall creates a copy of the server's certificate and signs it with the Forward Untrust certificate and sends it to the client. In this case, the client sees a block page warning that the site they're attempting to connect to is not trusted and the client can choose to proceed or terminate the session. When the client authenticates the certificate, the SSL session is established with the firewall functioning as a trusted forward proxy to the site that the client is accessing.

As the firewall continues to receive SSL traffic from the server that is destined for the client, it decrypts the SSL traffic into clear text traffic and applies security policies to the traffic. The traffic is then re-encrypted on the firewall and the firewall forwards the encrypted traffic to the client.

Figure: SSL Forward Proxy shows this process in detail.

 © Palo Alto Networks, Inc.

**Figure: SSL Forward Proxy**

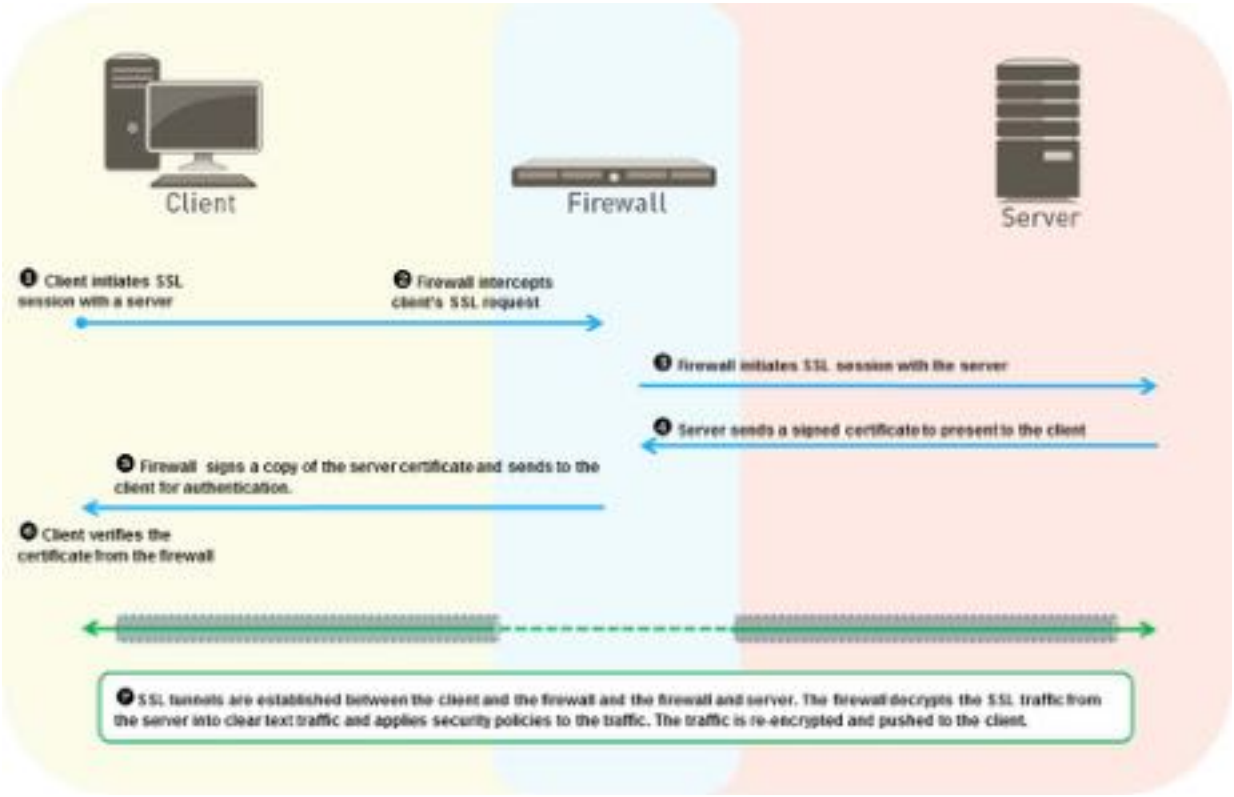


**See** Configure SSL Forward Proxy **for details on configuring SSL Forward Proxy.**

## SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL traffic from a client to a targeted server (any server you have the certificate for and can import it onto the firewall). For example, if an employee is remotely connected to a web server hosted on the company network and is attempting to add restricted internal documents to his Dropbox folder (which uses SSL for data transmission), SSL Inbound Inspection can be used to ensure that the sensitive data does not move outside the secure company network by blocking or restricting the session.

Configuring SSL Inbound Inspection includes importing the targeted server's certificate and key on to the firewall. Because the targeted server's certificate and key is imported on the firewall, the firewall is able to access the SSL session between the server and the client and decrypt and inspect traffic transparently, rather than functioning as a proxy. The firewall is able to apply security policies to the decrypted traffic, detecting malicious content and controlling applications running over this secure channel.

Figure: SSL Inbound Inspection shows this process in detail.

**Figure: SSL Inbound Inspection**

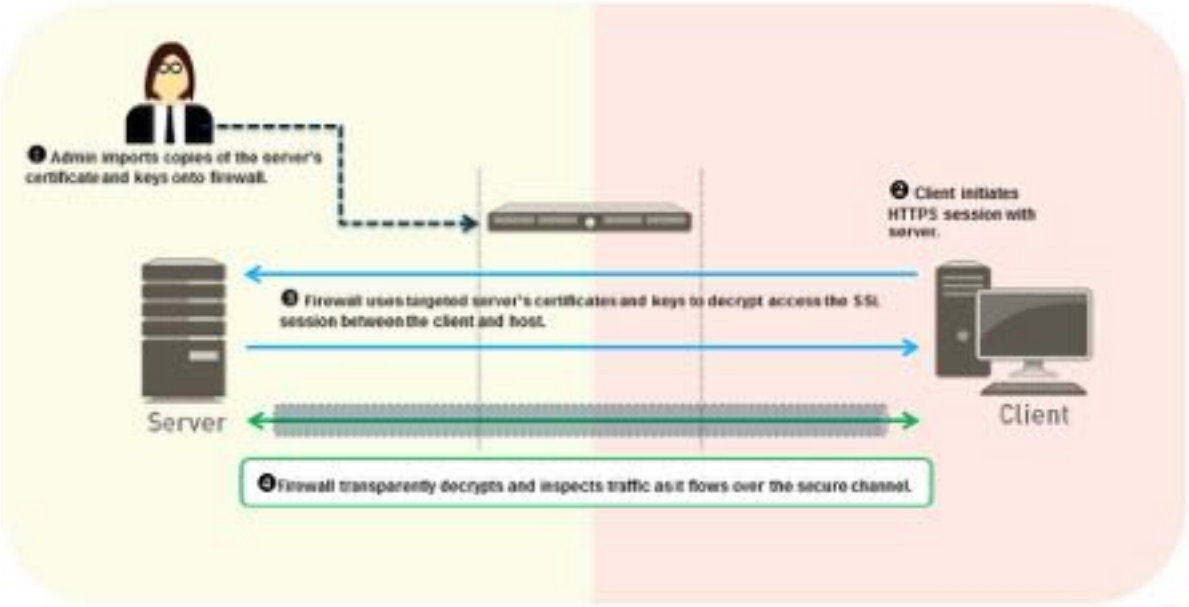


See Configure SSL Inbound Inspection for details on configuring SSL Inbound Inspection.

 © Palo Alto Networks, Inc.

## SSH Proxy

SSH Proxy provides the capability for the firewall to decrypt inbound and outbound SSH connections passing through the firewall, in order to ensure that SSH is not being used to tunnel unwanted applications and content. SSH decryption does not require any certificates and the key used for SSH decryption is automatically generated when the firewall boots up. During the boot up process, the firewall checks to see if there is an existing key. If not, a key is generated. This key is used for decrypting SSH sessions for all virtual systems configured on the device. The same key is also used for decrypting all SSH v2 sessions.

In an SSH Proxy configuration, the firewall resides between a client and a server. When the client sends an SSH request to the server, the firewall intercepts the request and forwards the SSH request to the server. The firewall then intercepts the server's response and forwards the response to the client, establishing an SSH tunnel between the firewall and the client and an SSH tunnel between the firewall and the server, with firewall functioning as a proxy. As traffic flows between the client and the server, the firewall is able to distinguish whether the SSH traffic is being routed normally or if it is using SSH tunneling (port forwarding). Content and threat inspections are not performed on SSH tunnels; however, if SSH tunnels are identified by the firewall, the SSH tunneled traffic is blocked and restricted according to configured security policies.

Figure: SSH Proxy Decryption shows this process in detail.

**Figure: SSH Proxy Decryption**

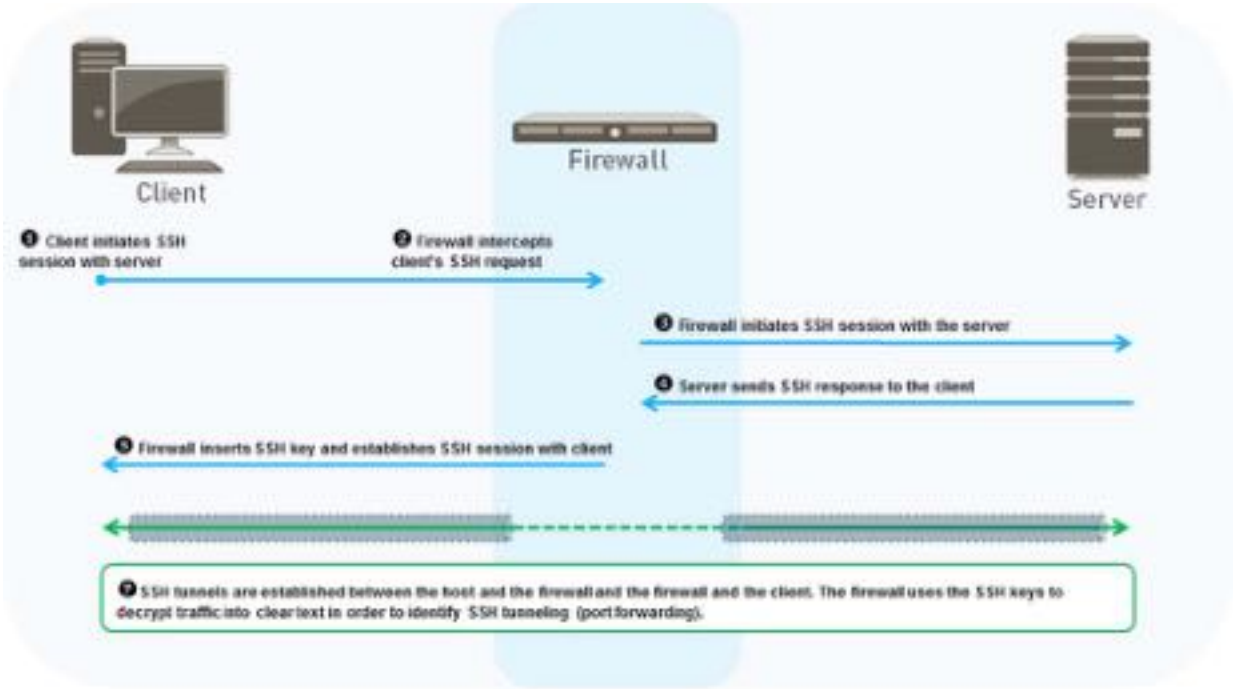


See Configure SSH Proxy for details on configuring an SSH Proxy policy.

## Decryption Exceptions

Traffic can also be excluded from decryption according to matching criteria (using a decryption policy), a targeted server's traffic can be excluded from decryption (using certificates), and some applications are excluded from decryption by default.

Applications that do not function properly when decrypted by the firewall and are automatically excluded from SSL decryption. The applications that are excluded from SSL decryption by default are excluded because these applications often fail when decrypted due to the application looking for specific details in the certificate that might not be present in the certificate generated for SSL Forward Proxy. Refer to the KB article List of Applications Excluded from SSL Decryption for a current list of applications excluded by default from SSL decryption on the firewall.

You can configure decryption exceptions for certain URL categories or applications that either do not work properly with decryption enabled or for any other reason, including for legal or privacy purposes. You can use a decryption policy to exclude traffic from decryption based on source, destination, and URL category. For example, with SSL decryption enabled, you can exclude traffic that is categorized as financial or health-related from decryption, using the URL category selection. To create a decryption policy that excludes traffic from decryption.

You can also exclude servers from SSL decryption based on the Common Name (CN) in the server's certificate. For example, if you have SSL decryption enabled but have certain servers that you do not want included in SSL decryption, such as the web services for your HR systems, you can exclude those servers from decryption by importing the server certificate onto the firewall and modifying the certificate to be an **SSL Exclude Certificate**.

To exclude traffic from decryption based on application, source, destination, URL category or to exclude a specific server's traffic from decryption, see Configure Decryption Exceptions.

 © Palo Alto Networks, Inc.

## Decryption Port Mirroring

The Decryption Port mirror feature provides the capability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures–such as NetWitness or Solera–for archiving and analysis. This feature is necessary for organizations that require comprehensive data capture for forensic and historical purposes or data leak prevention (DLP) functionality. Decryption port mirroring is available on PA-7050, PA-5000 Series and PA-3000 Series platforms only and requires that a free license be installed to enable this feature.

Keep in mind that the decryption, storage, inspection, and/or use of SSL traffic is governed in certain countries and user consent might be required in order to use the decryption port mirror feature. Additionally, use of this feature could enable malicious users with administrative access to the firewall to harvest usernames, passwords, social security numbers, credit card numbers, or other sensitive information submitted using an encrypted channel. Palo Alto Networks recommends that you consult with your corporate council before activating and using this feature in a production environment.

Figure: Decryption Port Mirroring shows the process for decryption port mirroring and the section Configure Decryption Port Mirroring describes how to license and use this feature.

**Figure: Decryption Port Mirroring**

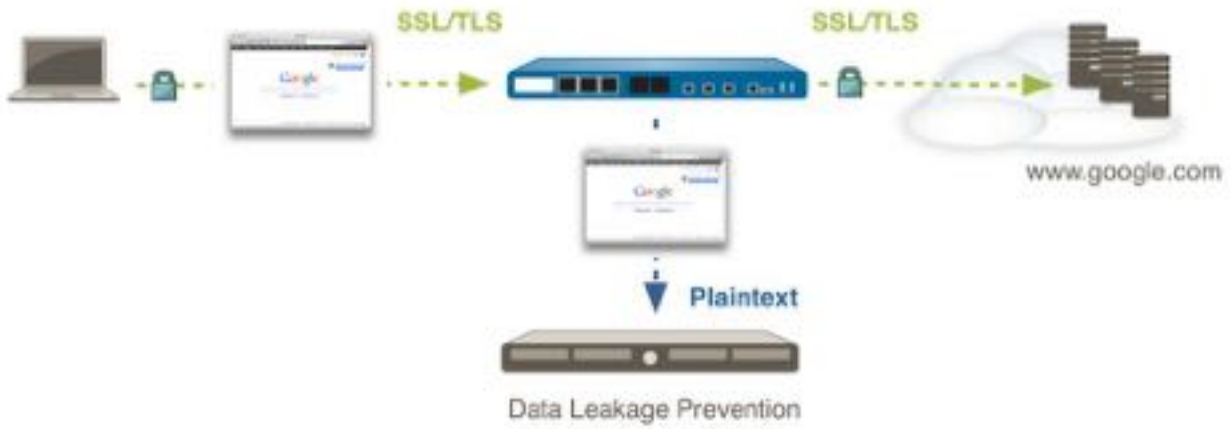

# Configure SSL Forward Proxy

Configuring SSL Forward Proxy decryption on the firewall requires setting up the certificates needed for SSL Forward Proxy decryption and creating an SSL Forward Proxy decryption policy. The firewall can use self-signed certificates or certificates signed by an enterprise CA to perform SSL Forward Proxy decryption.

Use the following task to configure SSL Forward Proxy, including how to set up the certificates and create a decryption policy.

| **Configure SSL Forward Proxy** | | |
|---|---|---|
| Step 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. | View configured interfaces on the **Network > Interfaces > Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is. |
| Step 2 | Configure the forward trust certificate. Use either a self-signed certificate or a certificate signed by an enterprise CA.<br>**Using self-signed certificates**<br>When the certificate of the server that the client is connecting to is signed by a CA that is on the firewall's trusted CA list, the firewall signs a copy of the server's certificate with a self-signed forward trust certificate to present to the client for authentication. In this case, the self-signed certificate must be imported onto each client system so that the client recognizes the firewall as a trusted CA.<br>Use self-signed certificates for SSL Forward Proxy decryption if you do not use an enterprise CA or if you are only intended to perform decryption for a limited number of client systems (or if you are planning to use a centralized deployment). | To use a self-signed certificate:<br>1. Select **Device > Certificate Management > Certificates**.<br>2. Click **Generate** at the bottom of the window.<br>3. Enter a **Certificate Name**, such as *my-fwd-trust*.<br>4. Enter a **Common Name**, such as **192.168.2.1**. This should be the IP or FQDN that will appear in the certificate. In this case, we are using the IP of the trust interface. Avoid using spaces in this field.<br>5. Leave the **Signed By** field blank.<br>6. Click the **Certificate Authority** check box to enable the firewall to issue the certificate. Selecting this check box creates a certificate authority (CA) on the firewall that is imported to the client browsers, so clients trust the firewall as a CA.<br>7. Click **Generate** to generate the certificate.<br>8. Click the new certificate *my-fwd-trust* to modify it and enable the **Forward Trust Certificate** option.<br>9. Export the forward trust certificate for import into client systems by highlighting the certificate and clicking **Export** at the bottom of the window. Choose PEM format, and do not select the **Export private key** option. Because the certificate is self-signed, import it into the browser trusted root CA list on the client systems in order for the clients to trust it. When importing to the client browser, ensure the certificate is added to the Trusted Root Certification Authorities certificate store. On Windows systems, the default import location is the Personal certificate store. You can also simplify this process by using a centralized deployment, such as an Active Directory Group Policy Object (GPO).<br> If the forward trust certificate is not imported on the client systems, users will see certificate warnings for each SSL site they visit.<br>10. Click **OK** to save. |


© Palo Alto Networks, Inc.

## Configure SSL Forward Proxy

### Using an Enterprise CA

An enterprise CA can issue a signing certificate which the firewall can use to then sign the certificates for sites requiring SSL decryption. Send a Certificate Signing Request (CSR) for the enterprise CA to sign and validate. The firewall can then use the signed enterprise CA certificate for SSL Forward Proxy decryption. Because the enterprise CA is already trusted by the client systems, with this option, you do not need to distribute the certificate to client systems prior to configuring decryption.

To use an enterprise CA signed certificate, generate a CSR:

1. Select **Device > Certificate Management > Certificates** and click **Generate**.
2. Enter a **Certificate Name**, such as *my-fwd-proxy*.
3. In the **Signed By** dropdown, select **External Authority (CSR)**.
4. (Optional) If your enterprise CA requires it, add **Certificate Attributes** to further identify the firewall details, such as Country or Department.
5. Click **OK** to save the CSR. The pending certificate is now displayed on the **Device Certificates** tab.
6. Export the CSR:
   a. Select the pending certificate displayed on the **Device Certificates** tab.
   b. Click **Export** to download and save the certificate file.

   Leave **Export private key** unselected in order to ensure that the private key remains securely on the firewall.
   c. Click **OK**.
   d. Provide the certificate file to your enterprise CA. When you receive the signed enterprise CA certificate from your enterprise CA, save the signed enterprise CA certificate for import onto the firewall.
7. Import the signed enterprise CA onto the firewall:
   a. Select **Device > Certificate Management > Certificates** and click **Import**.
   b. Enter the pending **Certificate Name** exactly (in this case, *my-fwd-trust*). The **Certificate Name** that you enter must exactly match the pending certificate's name in order for the pending certificate to be validated.
   c. Select the signed **Certificate File** that you received from your enterprise CA.
   d. Click **OK**. The certificate is displayed as valid with the Key and CA check boxes selected.
   e. Select the validated certificate, in this case, *my-fwd-proxy*, to enable it as a **Forward Trust Certificate** to be used for SSL Forward Proxy decryption.
   f. Click **OK**.

| Configure SSL Forward Proxy | | |
|---|---|---|
| Step 3 | Configure the forward untrust certificate.<br><br>With SSL Forward Proxy decryption, when the site the client is connecting to uses a certificate signed by a CA that is not in the firewall's trusted CA list, the firewall presents a forward untrust certificate to the client. The forward untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites with untrusted certificates. | 1. Click **Generate** at the bottom of the certificates page.<br>2. Enter a **Certificate Name**, such as *my-fwd-untrust*.<br>3. Set the **Common Name**, for example **192.168.2.1**. Leave **Signed By** blank.<br>4. Click the **Certificate Authority** check box to enable the firewall to issue the certificate.<br>5. Click **Generate** to generate the certificate.<br>6. Click **OK** to save.<br>7. Click the new *my-ssl-fw-untrust* certificate to modify it and enable the **Forward Untrust Certificate** option.<br> Do not export the forward untrust certificate for import into client systems. If the forward trust certificate is imported on client systems, the users will not see certificate warnings for SSL sites with untrusted certificates.<br>8. Click **OK** to save. |
| Step 4 | (Optional) Create a Decryption profile.<br><br>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. An SSL Forward Proxy decryption profile can be used to perform checks for server certificates, unsupported modes, and failures and block or restrict traffic accordingly. For a complete list of checks that can be performed, navigate to **Objects > Decryption Profiles** on the firewall and click the help icon. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Select the **SSL Forward Proxy** tab to block and control specific aspects of SSL tunneled traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting **Block sessions if resources not available**.<br>3. Click **OK t**o save the profile. |

 © Palo Alto Networks, Inc.

| Configure SSL Forward Proxy | |
|---|---|
| Step 5 Configure a decryption policy. | 1. Select **Policies > Decryption** and click **Add**.<br>2. On the **General** tab, give the policy a descriptive **Name**.<br>3. On the **Source** and **Destination** tabs, select **Any** for the **Source Zone** and **Destination Zone** to decrypt all SSL traffic destined for an external server. If you want to specify traffic from or to certain sources or destinations for decryption, click **Add**.<br>4. In the **URL Category** tab, leave **Any** to decrypt all traffic. If you only want to apply this profile to certain website categories, click **Add**.<br> Selecting a URL Category is useful when excluding certain sites from decryption. See Configure Decryption Exceptions.<br>5. On the **Options** tab, select **Decrypt** and select **SSL Forward Proxy** as the **Type** of decryption to perform.<br>6. (Optional) Select a **Decryption Profile** to apply additional settings to decrypted traffic (see Step 4).<br>7. Click **OK** to save. |
| Step 6 **Commit** the configuration. | With the an SSL Forward Proxy decryption policy enabled, all traffic identified by the policy is decrypted. Decrypted traffic is blocked and restricted according to the profiles configured on the firewall (including the decryption profiles associated with the policy and Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles). Traffic is re-encrypted as it exits the firewall. |

# Configure SSL Inbound Inspection

Configuring SSL Inbound Inspection includes installing the targeted server's certificate on the firewall and creating an SSL Inbound Inspection decryption policy.

Use the following task to configure SSL Inbound Inspection.

| Configure SSL Inbound Inspection | | |
|---|---|---|
| Step 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. | View configured interfaces on the **Network > Interfaces > Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is. |
| Step 2 | Ensure that the targeted server's certificate is installed on the firewall. | On the web interface, select **Device > Certificate Management > Certificates > Device Certificates** to view certificates installed on the firewall.<br>To import the targeted server's certificate onto the firewall:<br>1. On the **Device Certificates** tab, select **Import**.<br>2. Enter a descriptive **Certificate Name**.<br>3. Browse for and select the targeted server's **Certificate File**.<br>4. Click **OK**. |
| Step 3 | (Optional) Create a Decryption profile.<br>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. An SSL Inbound Inspection decryption profile can be used to perform checks for unsupported modes and failures and block or restrict traffic accordingly. For a complete list of checks that can be performed, select **Objects > Decryption Profiles** and then click the help icon. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Select the **SSL Inbound Inspection** tab to block and control specific aspects of SSL traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting **Block sessions if resources not available**.<br>3. Click **OK t**o save the profile. |


© Palo Alto Networks, Inc.

| Configure SSL Inbound Inspection | |
|---|---|
| Step 4 Configure a decryption policy. | 1. Select **Policies > Decryption** and click **Add .**<br>2. On the **General** tab, give the policy a descriptive **Name**.<br>3. On the **Destination** tab, **Add** the **Destination Address** of the targeted server.<br>4. In the **URL Category** tab, leave **Any** to decrypt all traffic. If you only want to apply this profile to certain website categories, click **Add**.<br> Selecting a URL Category is useful when excluding certain sites from decryption. See Configure Decryption Exceptions.<br>5. On the **Options** tab, select **Decrypt** and select **SSL Inbound Inspection** as the **Type** of traffic to decrypt.<br>Select the **Certificate** for the internal server that is the destination of the inbound SSL traffic.<br>6. (Optional) Select a **Decryption Profile** to apply additional settings to decrypted traffic.<br>7. Click **OK** to save. |
| Step 5 **Commit** the configuration. | With an SSL Inbound Inspection decryption policy enabled, all SSL traffic identified by the policy is decrypted and inspected. Decrypted traffic is blocked and restricted according to the profiles configured on the firewall (including the decryption profiles associated with the policy and Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles). Traffic is re-encrypted as it exits the firewall. |

# Configure SSH Proxy

Configuring SSH Proxy does not require certificates and the key used to decrypt SSH sessions is generated automatically on the firewall during boot up.

Use the following task to configure SSH Proxy decryption.

| Configure SSH Proxy Decryption | | |
|---|---|---|
| Step 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. Decryption can only be performed on virtual wire, Layer 2, or Layer 3 interfaces. | View configured interfaces on the **Network > Interfaces > Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is. |
| Step 2 | (Optional) Create a Decryption profile.<br>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. The decryption profile can be used to perform checks for server certificates, unsupported modes, and failures and block or restrict traffic accordingly. For a complete list of checks that can be performed, navigate to **Objects > Decryption Profiles** on the firewall and then click the help icon. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Select the **SSH** tab to block and control specific aspects of SSH tunneled traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting **Block sessions if resources not available**.<br>3. Click **OK** to save the profile. |
| Step 3 | Configure a decryption policy. | 1. Select **Policies > Decryption** and click **Add**.<br>2. On the **General** tab, give the policy a descriptive **Name**.<br>3. On the **Source** and **Destination** tabs, select **Any** to decrypt all SSH traffic.<br>4. On the **URL Category** tab, select **Any** to decrypt all SSH traffic.<br>5. On the **Options** tab, select **Decrypt** and select **SSH Proxy** as the **Type** of traffic to decrypt.<br>6. (Optional) Select a **Decryption Profile** to apply additional settings to decrypted traffic.<br>7. Click **OK** to save. |
| Step 4 | **Commit** the configuration. | With the an SSH Proxy decryption policy enabled, all SSH traffic identified by the policy is decrypted and identified as either regular SSH traffic or as SSH tunneled traffic. SSH tunneled traffic is blocked and restricted according to the profiles configured on the firewall. Traffic is re-encrypted as it exits the firewall. |


© Palo Alto Networks, Inc.

# Configure Decryption Exceptions

You can purposefully exclude traffic from decryption based on match criteria, such as the application, the traffic's source or destination, or the URL category. You can also exclude a specific server's traffic from decryption. See the following topics to configure Decryption Exceptions:

- Exclude Traffic From Decryption
- Exclude a Server From Decryption

## Exclude Traffic From Decryption

To purposefully exclude applications or certain traffic from other existing SSL or SSH decryption policies, you can create a new decryption policy that defines the traffic to exclude from decryption and with the **No Decrypt** action selected in the policy. You can define traffic for policy-based exclusion according to match criteria, such as application, source, destination, or URL categories. Make sure the decryption policy that excludes traffic from decryption is listed first in your decryption policy list by dragging and dropping the policy above the other decryption policies.

See the following procedure to configure a decryption policy that excludes traffic from SSL or SSH decryption.

| **Exclude Traffic from a Decryption Policy** | | |
|---|---|---|
| Step 1 | Create a decryption policy.<br><br>Use a decryption policy to exclude traffic from decryption according to the traffic's source and destination zones or addresses and URL categories. This example shows how to exclude traffic categorized as financial or health-related from SSL Forward Proxy decryption. | 1. Go to **Policies > Decryption** and click **Add**.<br>2. Give the policy a descriptive **Name**, such as *No-Decrypt-Finance-Health*.<br>3. On the **Source** and **Destination** tabs, select **Any** for the **Source Zone** and **Destination Zone** to apply the *No-Decrypt-Finance-Health* rule to all SSL traffic destined for an external server.<br>4. On the **URL Category** tab, **Add** the URL categories financial-services and health-and-medicine to the policy, specifying that traffic that matches these categories will not be decrypted.<br>5. On the **Options** tab, select **No Decrypt** and select the **Type** of decryption policy you are excluding the traffic from. For example, to exclude traffic categorized as financial or health-related from a separately configured SSL Forward Proxy decryption policy, select **SSL Forward Proxy** as the **Type**.<br>6. Click **OK** to save the *No-Decrypt-Finance-Health* decryption policy. |
| Step 2 | Move the decryption policy to the top of the list of decryption policies. | On the **Decryption > Policies** page, select the policy *No-Decrypt-Finance-Health*, and click **Move Up** until it appears at the top of the list (or you can drag and drop).<br><br>The order in which the decryption policies are listed is the order in which they are applied to network traffic. Moving the policy with the **No Decrypt** action applied to the top of the list ensures that the specified traffic is not decrypted according to another configured policy. |

| Exclude Traffic from a Decryption Policy | |
|---|---|
| Step 3 **Commit** the configuration. | A decryption policy with **No Decrypt** enabled ensures that the specified traffic is remains encrypted as it flows through the firewall, and that the traffic is not decrypted according to other decryption policies configured and listed on the **Policies > Decryption** page. |

## Exclude a Server From Decryption

You can exclude a targeted server's traffic from SSL decryption based on the Common Name (CN) in the server's certificate. For example, if you have SSL decryption enabled, you could configure a decryption exception for server on your corporate network that hosts the web services for your HR systems. See the following procedure to configure modify a server's certificate so that the targeted server's traffic is excluded from decryption:

| Exclude a Server from Decryption | |
|---|---|
| Step 1 | Import the targeted server's certificate onto the firewall:<br>1. On the **Device > Certificate Management > Certificates > Device Certificates** tab, select **Import**.<br>2. Enter a descriptive **Certificate Name**.<br>3. Browse for and select the targeted server's **Certificate File**.<br>4. Click **OK**. |
| Step 2 | Select the targeted server's certificate on the **Device Certificates** tab and enable it as an SSL Exclude Certificate.<br>With the targeted server's certificate imported on the firewall and designated as an SSL Exclude Certificate, the server's traffic is not decrypted as it passes through the firewall. |

 © Palo Alto Networks, Inc.

# Configure Decryption Port Mirroring

Before you can enable decryption port mirroring, you must obtain and install a Decryption Port Mirror license. The license is free of charge and can be activated through the support portal as described in the following procedure. After you install the Decryption Port Mirror license and reboot the firewall, you can enable decryption port mirroring. Enabling decryption port mirroring includes enabling the forwarding of decrypted traffic and configuring a decrypt mirror interface. You can then create a decryption profile that specifies the interface and attach it to a decryption policy. To learn more about implementing Decryption Port Mirroring, see Decryption Port Mirroring.

Use the following procedure to obtain and install a Decryption Port Mirror license and configure Decryption Port Mirroring.

| **Configure Decryption Port Mirroring** | |
|---|---|
| Step 1 Request a license for each device on which you want to enable decryption port mirroring. | 1. Log in to the Palo Alto Networks Support site and navigate to the **Assets** tab.<br>2. Select the device entry for the device you want to license and select **Actions**.<br>3. Select **Decryption Port Mirror**. A legal notice displays.<br>4. If you are clear about the potential legal implications and requirements, click **I understand and wish to proceed**.<br>5. Click **Activate**.<br> 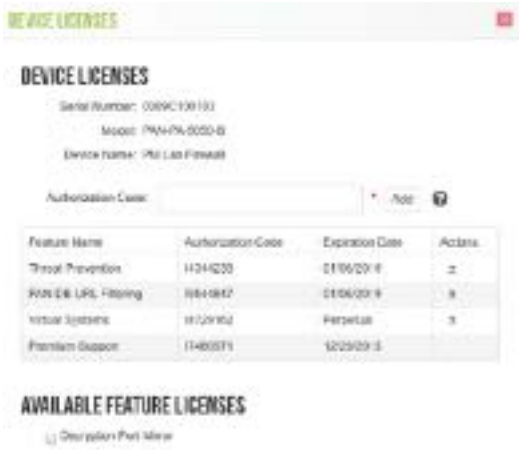 |
| Step 2 Install the Decryption Port Mirror license on firewall. | 1. From the firewall's web interface, select **Device > Licenses**.<br>2. Click **Retrieve license keys from license server**.<br>3. Verify that the license has been activated on the firewall.<br> <br>4. Reboot the firewall (**Device > Setup > Operations**). This feature will not be available for configuration until PAN-OS reloads. |

| Configure Decryption Port Mirroring (Continued) | | |
|---|---|---|
| Step 3 | Enable the ability to mirror decrypted traffic. Superuser permission is required to perform this step. | **On a firewall with a single virtual system:**<br>1. Select **Device > Setup > Content - ID**.<br>2. Select the **Allow forwarding of decrypted content** check box.<br>3. Click **OK** to save.<br>**On a firewall with multiple virtual systems:**<br>1. Select **Device > Virtual System.**<br>2. Select a Virtual System to edit or create a new Virtual System by selecting **Add**.<br>3. Select the **Allow forwarding of decrypted content** check box.<br>4. Click **OK** to save. |
| Step 4 | Configure a decrypt mirror interface. | 1. Select **Network > Interfaces > Ethernet**.<br>2. Select the Ethernet interface that you want to configure for decryption port mirroring.<br>3. Select **Decrypt Mirror** as the **Interface Type**.<br>This interface type will only appear if the Decryption Port Mirror license is installed.<br>4. Click **OK** to save. |
| Step 5 | Configure a Decryption Profile to enable decryption port mirroring. | 1. Select **Objects > Decryption Profile**.<br>2. Select the **Interface** to use for **Decryption Mirroring**.<br>The **Interface** drop-down contains all Ethernet interfaces that have been defined as the type: **Decrypt Mirror**.<br>3. Specify whether to mirror decrypted traffic before or after policy enforcement.<br>By default, the firewall will mirror all decrypted traffic to the interface before security policies lookup, which allows you to replay events and analyze traffic that generates a threat or triggers a drop action. If you want to only mirror decrypted traffic after security policy enforcement, select the **Forwarded Only** check box. With this option, only traffic that is forwarded through the firewall is mirrored. This option is useful if you are forwarding the decrypted traffic to other threat detection devices, such as a DLP device or another intrusion prevention system (IPS).<br>4. Click **OK** to save the decryption profile. |
| Step 6 | Set a decryption policy for decryption port mirroring. | 1. Select **Policies > Decryption**.<br>2. Click **Add** to configure a decryption policy or select an existing decryption policy to edit.<br>3. In the **Options** tab, select **Decrypt** and the **Decryption Profile** created in Step 4.<br>4. Click **OK** to save the policy. |
| Step 7 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# URL Filtering

The Palo Alto Networks URL filtering solution is a powerful PAN-OS feature that is used to monitor and control how users access the web over HTTP and HTTPS. The following topics provide an overview of URL filtering, configuration and troubleshooting information, and best practices for getting the most out of this feature:

- URL Filtering Overview
- URL Filtering Concepts
- PAN-DB Categorization Workflow
- Configure URL Filtering
- Enable Safe Search Enforcement
- URL Filtering Use Case Examples
- Troubleshoot URL Filtering

# URL Filtering Overview

The Palo Alto Networks URL filtering feature compliments the App-ID feature by enabling you to configure your firewall to identify and control access to web (HTTP and HTTPS) traffic. By implementing URL filtering profiles in security policies and by using URL categories as a match criteria in policies (captive portal, decryption, security, and QoS), you will gain complete visibility and control of the traffic that traverses your firewall and will be able to safely enable and control how your users access the web.

The Palo Alto Networks URL filtering solution utilizes a URL filtering database that contains millions of websites and each website is placed in one of approximately 60 different categories. A URL filtering profile that contains the list of categories is then applied to a security policy that allows web traffic (HTTP/HTTPS) from the internal users to the Internet. After the URL filtering profile is applied and the alert or block action is set on a category, you will gain complete visibility into the websites that users access and can then decide which websites or website categories should be allowed, blocked, or logged. You can also define a list of URLs in the URL filtering profile that will always be blocked or allowed and you can create custom URL categories that contains a list of URLs that can be used the same way as the default category list. These same URL categories can also be used as a match criteria in other policies, such as captive portal, decryption, and QoS.

## URL Filtering Vendors

Palo Alto Networks firewalls support two vendors for URL filtering purposes:

- **PAN-DB**—A Palo Alto Networks developed URL filtering database that is tightly integrated into PAN-OS by utilizing high-performance local caching to perform maximum inline performance for URL lookups while a distributed cloud architecture provides coverage for the latest websites. To view a list of PAN-DB URL filtering categories, refer to https://urlfiltering.paloaltonetworks.com/CategoryList.aspx.
- **BrightCloud**—A third-party URL database that is owned by Webroot, Inc. and is integrated into PAN-OS firewalls. For information on the BrightCloud URL database, visit http://brightcloud.com.

## Interaction Between App-ID and URL Categories

The Palo Alto Networks URL filtering feature in combination with Application Identification (App-ID) provides unprecedented protection against a full spectrum of legal, regulatory, productivity, and resource utilization risks. While App-ID gives you control over what applications users can access, URL filtering provides control over related web activity. When combined with User-ID, you can also apply these controls based on users and groups.

With today's application landscape and the way many applications use HTTP and HTTPS, you will need to determine when to use App-ID and when to use URL filtering in order to define comprehensive web access policies. In most cases, if an App-ID signature exists, you will want to use App-ID to control the content at the application level. For example, although you can control access to Facebook and/or LinkedIn using URL filtering, this would not block the use of all related applications, such as email, chat, as well as any new application that is introduced after you implement your policy.

 © Palo Alto Networks, Inc.

In some cases, you will want to use both URL filtering and App-ID, but to ensure that conflicts do not occur in your policies, it is important to understand how these features work together. Palo Alto Networks generates signatures for many applications and those signatures can be very granular in regard to various features within the web-based applications—whereas URL filtering would only apply actions based on a specific website or URL category. For example, you may want to block social networking sites in general, but want to allow a few sites in that category to be accessible to specific departments and then control what features of the website are available to the allowed users. For more information, see URL Filtering Use Case Examples.

© Palo Alto Networks, Inc.

# URL Filtering Concepts

The following topics describe the URL filtering components and how they are used on a Palo Alto Networks firewall:


- URL Categories
- URL Filtering Profiles
- URL Category as Policy Match Criteria


## URL Categories

Each website defined in the URL filtering database is assigned one of approximately 60 different categories. These categories can then be used in a URL filtering profile to block or allow access based on category, or you can configure the firewall to use a category as a match criteria in policy. For example, to block all gaming websites, in the URL filtering profile you would set the block action for the URL category *games*. As an example of using a URL category as a match criteria in a policy, you could use the URL category *streaming-media* in a QoS policy to apply bandwidth controls to all websites that are categorized as streaming media.

By grouping websites into categories, it makes it easy to define actions based on certain types of websites. In addition to the standard URL categories, there are three additional categories:

- **Not-resolved**—Indicates that the website was not found in the local URL filtering database and the firewall was unable to connect to the cloud database to check the category. When a URL category lookup is performed, the firewall first checks the dataplane cache for the URL, if no match is found, it will then check the management plane cache, and if no match is found there, it queries the URL database in the cloud.

  When deciding on what action to take for traffic that is categorized as *not-resolved*, be aware that setting the action to block may be very disruptive to users.

  For more information on troubleshooting lookup issues, see Troubleshoot URL Filtering.

- **Private-ip-addresses**—Indicates that the website is a single domain (no sub-domains), the IP address is in the private IP range, or the URL root domain is unknown to the cloud.

- **Unknown**—The website has not yet been categorized, so it does not exist in the URL filtering database on the firewall or in the URL cloud database.

  When deciding on what action to take for traffic categorized as *unknown*, be aware that setting the action to block may be very disruptive to users because there could be a lot of valid sites that are not in the URL database yet. If you do want a very strict policy, you could block this category, so websites that do not exist in the URL database cannot be accessed.

 © Palo Alto Networks, Inc.

# URL Filtering Profiles

A URL filtering profile is a collection of URL filtering controls that are applied to individual security policies to enforce your web access policy. The firewall comes with a default profile that is configured to block websites such as known malware sites, phishing sites, and adult content sites. You can use the default profile in a security policy, clone it to be used as a starting point for new URL filtering profiles, or add a new URL profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories. For example, you may want to block social-networking sites, but allow some websites that are part of the social-networking category.

The section describes how URL filtering profiles are applied and the various options that can be defined:

- URL Filtering Actions
- Block and Allow Lists
- Safe Search Enforcement
- Container Pages

## URL Filtering Actions

Each URL filtering category can be set to perform the following actions:

| Action | Description |
|---|---|
| Alert | The website is allowed and a log entry is generated in the URL filtering log. |
| Allow | The website is allowed and no log entry is generated. |
| Block | The website is blocked and the user will see a response page and will not be able to continue to the website. A log entry is generated in the URL filtering log. |
| Continue | The user will be prompted with a response page indicating that the site has been blocked due to company policy, but the user is prompted with the option to continue to the website. The continue action is typically used for categories that are considered benign and is used to improve the user experience by giving them the option to continue if they feel the site is incorrectly categorized. The response page message can be customized to contain details specific to your company. A log entry is generated in the URL filtering log.<br> The *Continue* page will not be displayed properly on client machines that are configured to use a proxy server. |
| Override | The user will see a response page indicating that a password is required to allow access to websites in the given category. With this option, the security admin or helpdesk person would provide a password that will grant temporary access to all websites in the given category. A log entry is generated in the URL filtering log.<br> The *Override* page will not be displayed properly on client machines that are configured to use a proxy server. |

| Action | Description |
|---|---|
| None | The *None* action only applies to custom URL categories. The purpose of selecting *None* is to ensure that if multiple URL profiles exist, the custom category will not have any impact on other profiles. For example, if you have two URL profiles and the custom URL category is set to block in one of the profiles, the other profile should have the action set to *None* if you do not want it to apply.<br>Also, in order to delete a custom URL category, it must be set to none in any profile where it is used. |

## Block and Allow Lists

Block and allow lists allow you to define specific URLs or IP addresses in the URL filtering profile that are always allowed or always blocked, regardless of the action defined for the URL category. When entering URLs in the *Block List* or *Allow List*, enter each URL or IP address in a new row separated by a new line. When using wildcards in the URLs, follow these rules:

- Do not include HTTP and HTTPS when defining URLs. For example, enter www.paloaltonetworks.com or paloaltonetworks.com instead of https://www.paloaltonetworks.com.
- Entries in the block list must be an exact match and are case-insensitive.

  For example: If you want to prevent a user from accessing any website within the domain paloaltonetworks.com, you would also add *.paloaltonetworks.com, so whatever domain prefix (http://, www, or a sub-domain prefix such as mail.paloaltonetworks.com) is added to the address, the specified action will be taken. The same applies to the sub-domain suffix; if you want to block paloaltonetworks.com/en/US, you would need to add paloaltonetworks.com/* as well. Block and allow lists support wildcard patterns. The following characters are considered separators:

  .

  /

  ?

  &

  =

  ;

  +

  Every substring that is separated by the characters listed above is considered a token. A token can be any number of ASCII characters that does not contain any separator character or *. For example, the following patterns are valid:

  *.yahoo.com (tokens are: "*", "yahoo" and "com")

  www.*.com (tokens are: "www", "*" and "com")

  www.yahoo.com/search=* (tokens are: "www", "yahoo", "com", "search", "*")

  The following patterns are invalid because the character “*” is not the only character in the token.

  ww*.yahoo.com and www.y*.com

 © Palo Alto Networks, Inc.

## Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos in search query return traffic. On the firewall, you can Enable Safe Search Enforcement so that the firewall will block search results if the end user is not using the strictest safe search settings in the search query. The firewall can enforce safe search for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. This is a best-effort setting and is not guaranteed by the search providers to work with every website.

To use this feature you must enable the **Safe Search Enforcement** option in a URL filtering profile and attach it to a security policy. The firewall will then block any matching search query return traffic that is not using the strictest safe search settings. There are two methods for blocking the search results:

- Block Search Results that are not Using Strict Safe Search Settings—When an end user attempts to perform a search without first enabling the strictest safe search settings, the firewall blocks the search query results and displays the URL Filtering Safe Search Block Page. By default, this page will provide a URL to the search provider settings for configuring safe search.
- Enable Transparent Safe Search Enforcement—When an end user attempts to perform a search without first enabling the strict safe search settings, the firewall blocks the search results with an HTTP 503 status code and redirects the search query to a URL that includes the safe search parameters. You enable this functionality by importing a new URL Filtering Safe Search Block Page containing the Javascript for rewriting the search URL to include the strict safe search parameters. In this configuration, users will not see the block page, but will instead be automatically redirected to a search query that enforces the strictest safe search options. This safe search enforcement method requires Content Release version 475 or later and is only supported for Google, Yahoo, and Bing searches.

Also, because most search providers now use SSL to return search results, you must also configure a Decryption policy for the search traffic to enable the firewall to inspect the search traffic and enforce safe search.

Safe search enforcement enhancements and support for new search providers is periodically added in content releases. This information is detailed in the Application and Threat Content Release Notes. How sites are judged to be safe or unsafe is performed by each search provider, not by Palo Alto Networks.

Safe search settings differ by search provider as detailed in Table: Search Provider Safe Search Settings.

**Table: Search Provider Safe Search Settings**

| Search Provider | Safe Search Setting Description |
|---|---|
| Google/YouTube | Offers safe search on individual computers or network-wide through Google's safe search virtual IP address:<br><br>**Safe Search Enforcement for Google Searches on Individual Computers**<br><br>In the Google Search Settings, the **Filter explicit results** setting enables safe search functionality. When enabled, the setting is stored in a browser cookie as `FF=` and passed to the server each time the user performs a Google search.<br><br>Appending `safe=active` to a Google search query URL also enables the strictest safe search settings.<br><br>**Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address**<br><br>Google provides servers that Lock SafeSearch (forcesafesearch.google.com) settings in every Google and YouTube search. By adding a DNS entry for `www.google.com` and `www.youtube.com` (and other relevant Google and YouTube country subdomains) that includes a CNAME record pointing to `forcesafesearch.google.com` to your DNS server configuration, you can ensure that all users on your network are using strict safe search settings every time they perform a Google or YouTube search. Keep in mind, however, that this solution is not compatible with Safe Search Enforcement on the firewall. Therefore, if you are using this option to force safe search on Google, the best practice is to block access to other search engines on the firewall by creating custom URL categories and adding them to the block list in the URL filtering profile.<br><br> If you plan to use the Google Lock SafeSearch solution, consider configuring DNS Proxy (**Network > DNS Proxy**) and setting the inheritance source as the Layer 3 interface on which the firewall receives DNS settings from service provider via DHCP. You would configure the DNS proxy with **Static Entries** for www.google.com and www.youtube.com, using the local IP address for the forcesafesearch.google.com server. |
| Yahoo | Offers safe search on individual computers only. The Yahoo Search Preferences includes three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `vm=` and passed to the server each time the user performs a Yahoo search.<br><br>Appending `vm=r` to a Yahoo search query URL also enables the strictest safe search settings.<br><br> When performing a search on Yahoo Japan (yahoo.co.jp) while logged into a Yahoo account, end users must also enable the **SafeSearch Lock** option. |
| Bing | Offers safe search on individual computers or through their Bing in the Classroom program. The Bing Settings include three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `adlt=` and passed to the server each time the user performs a Bing search.<br><br>Appending `adlt=strict` to a Bing search query URL also enables the strictest safe search settings.<br><br>The Bing SSL search engine does not enforce the safe search URL parameters and you should therefore should consider blocking Bing over SSL for full safe search enforcement. |

 © Palo Alto Networks, Inc.

## Container Pages

A container page is the main page that a user accesses when visiting a website, but additional websites may be loaded within the main page. If the **Log Container page only** option is enabled in the URL filtering profile, only the main container page will be logged, not subsequent pages that may be loaded within the container page. Because URL filtering can potentially generate a lot of log entries, you may want to turn on this option, so log entries will only contain those URIs where the requested page file name matches the specific mime-types. The default set includes the following mime-types:

- application/pdf
- application/soap+xml
- application/xhtml+xml
- text/html
- text/plain
- text/xml

If you have enabled the **Log container page** only option, there may not always be a correlated URL log entry for threats detected by Antivirus or Vulnerability Protection.

# URL Category as Policy Match Criteria

URL categories can be used as a match criteria in a policy to provide more granularity in the policy. For example, you may have a decryption policy defined, but you would like specific websites to bypass decryption. To do this, you would configure a decryption policy with the no-decrypt action and a URL category would be defined as match criteria for the policy rule, so the policy would only match traffic flows to websites that are part of the specified category.

The following table describes the policy types that can utilize URL categories:

| Policy Type | Description |
|---|---|
| Captive Portal | To ensure that users authenticate before being allowed access to a specific category, you can attach a URL category as a match criteria for the captive portal policy. |
| Decryption | Decryption policies can use URL categories as a match criteria to determine if specified websites should be decrypted or not. For example, if you have a decryption policy with the action decrypt for all traffic between two zones, there may be specific website categories, such as *financial-services* and/or *health-and-medicine*, that should not be decrypted. In this case, you would create a new decryption policy with the action of *no-decrypt* that precedes the decrypt policy and then defines a list of URL categories as match criteria for the policy. By doing this, each URL category that is part of the no-decrypt policy will not be decrypted. You could also configure a custom URL category to define your own list of URLs that can then be used in the no decrypt policy. |

| Policy Type | Description |
|---|---|
| QoS | A QoS policy can use URL categories to determine throughput levels for specific website categories. For example, you may want to allow the streaming-media category, but limit throughput by adding the URL category as match criteria to the QoS policy. |
| Security | In security policies you can use URL categories both as a match criteria in the **Service/URL Category** tab, and in URL filtering profiles that are attached in the **Actions** tab.<br>If for example, the IT-security group in your company needs access to the hacking category, while all other users are denied access to the category, you must create the following rules:<br>• A security rule that allows the IT-Security group to access content categorized as hacking. The security rule references the *hacking* category in the S**ervices/URL Category** tab and IT-Security group in the **Users** tab.<br>• Another security rule that allows general web access for all users. To this rule you attach a URL filtering profile that blocks the hacking category.<br>The policy that allows access to hacking must be listed before the policy that blocks hacking. This is because security policy rules are evaluated top down, so when a user who is part of the security group attempts to access a hacking site, the policy rule that allows access is evaluated first and will allow the user access to the hacking sites. Users from all other groups are evaluated against the general web access rule which blocks access to the hacking sites. |

 © Palo Alto Networks, Inc.

# PAN-DB Categorization Workflow

This section describes the PAN-DB components and describes the URL categorization resolution workflow that occurs as users access various URLs through the firewall.

▲ PAN-DB URL Categorization Components
▲ PAN-DB URL Categorization Workflow

## PAN-DB URL Categorization Components

The following table describes the PAN-DB components in detail. The BrightCloud system works similarly, but does not use an initial seed database.

| Component | Description |
|---|---|
| URL Filtering Seed Database | The initial seed database downloaded to the firewall is a small subset of the database that is maintained on the Palo Alto Networks URL cloud servers. The reason this is done is because the full database contains millions of URLs and many of these URLs may never be accessed by your users. When downloading the initial seed database a region is selected (North America, Europe, APAC, Japan) and each region contains a subset of URLs most accessed for the given region. By doing this, the firewall will store a much smaller URL database, which greatly improves lookup performance. If a user accesses a website that is not in the local URL database, the full cloud database is queried and the firewall will then add the new URL to the local database. In other words, the local database on the firewall will be continually populated/customized based on user activity. Note that the local customized URL database will be cleared if the PAN-DB seed database is re-downloaded or if you change the URL database vendor from PAN-DB to BrightCloud. |
| Cloud Service | The PAN-DB cloud service is implemented using Amazon Web Services (AWS). AWS provides a distributed high performance and stable environment for seed database downloads and URL lookups for Palo Alto Networks firewalls and communication is performed over SSL. The AWS cloud systems hold the entire PAN-DB and is updated as new URLs are identified. The PAN-DB cloud service supports an automated mechanism to update the firewall's local URL database if the version does not match. Each time the firewall queries the cloud servers for URL lookups, it will also check for critical updates. If there have been no queries to the cloud servers for more than 30 minutes, the firewall will check for updates on the cloud systems.<br>The cloud system also provides a mechanism to submit URL category change requests. This is performed through the test-a-site service and is available directly from the device (URL filtering profile setup) and from the Palo Alto Networks Test A Site website. You can also submit URL categorization change request directly from URL filtering log on the firewall in the log details section. |

| Component | Description |
|---|---|
| Management Plane (MP) URL Cache | When PAN-DB is activated on the firewall, a seed database is downloaded from one of the PAN-DB cloud servers to initially populate the local cache, which is done to improve lookup performance. Each regional seed database contains the top URLs for the region and the size of the seed database (number of URL entries) also depends on the device platform. The URL MP cache is automatically written to the firewall's local drive every eight hours, before the firewall is rebooted, or when the cloud upgrades the URL database version on the firewall. After rebooting the firewall, the file that was saved to the local drive will be loaded to the MP cache. A least recently used (LRU) mechanism is also implemented in the URL MP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs. |
| Dataplane (DP) URL Cache | A subset of the MP cache and is a customized dynamic URL database that is stored in the dataplane (DP) and is used to improve URL lookup performance. The URL DP cache is cleared at each firewall reboot. The number of URLs that are stored in the URL DP cache varies by hardware platform and the current URLs stored in the TRIE (data structure). A least recently used (LRU) mechanism is implemented in the DP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs. Entries in the URL DP cache expire after a specified period of time and the expiration period cannot be changed by the administrator. |

## PAN-DB URL Categorization Workflow

When a user attempts to access a URL and the URL category needs to be determined, the firewall will compare the URL with the following components (in order) until a match has been found:

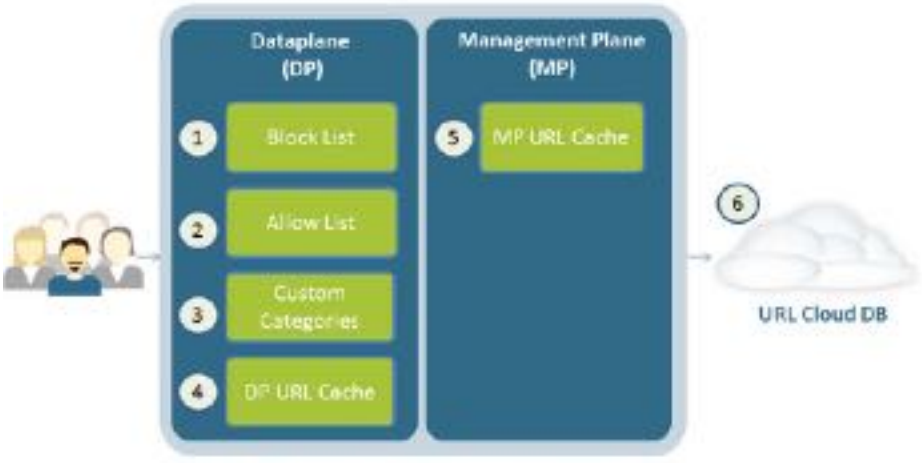

URL Categorization Resolution Process Flow

If a URL query matches an expired entry in the URL DP cache, the cache responds with the expired category, but also sends a URL categorization query to the management plane. This is done to avoid unnecessary delays in the DP, assuming that the frequency of changing categories is low. Similarly, in the URL MP cache, if a URL

 © Palo Alto Networks, Inc.

query from the DP matches an expired entry in the MP, the MP responds to the DP with the expired category and will also send a URL categorization request to the cloud service. Upon getting the response from the cloud, the firewall will resend the updated response to the DP.

As new URLs and categories are defined or if critical updates are needed, the cloud database will be updated. Each time the firewall queries the cloud for a URL lookup or if no cloud lookups have occurred for 30 minute, the database versions on the firewall be compared and if they do not match, an incremental update will be performed.

© Palo Alto Networks, Inc.

# Configure URL Filtering

This section describes the steps required to start using the URL filtering feature. After configuring URL filtering, you can monitor web activity and then determine what actions to take on specific websites and website categories. To control HTTPS traffic, the firewall must have a decryption policy in place between the zones that allow web traffic.

- Enable URL Filtering
- Determine URL Filtering Policy Requirements
- Define Website Controls


© Palo Alto Networks, Inc.

## Enable URL Filtering

To license and enable URL filtering on a Palo Alto Networks firewall:

| Enable URL Filtering | | |
|---|---|---|
| Step 1 | Obtain and install a URL filtering license and confirm that it is installed.<br> <br>The way PAN-DB and BrightCloud function after the URL filtering license expires is different. BrightCloud has an option in the URL profile to either allow all categories or block all categories if the license expires. With PAN-DB, if the license expires, URL filtering will continue to work based on the URL category information that exists in the dataplane and management plane caches, but URL cloud lookups and other cloud-based updates will not function until a valid license is installed. | 1. From **Device > Licenses** in the **License Management** section, select the license install method based on the type of license key you received. This will either be a key that you will retrieve from the license server, an authorization code, or a license file that is manually uploaded.<br>2. After the license is installed, confirm that a valid date is displayed in the **Date Expires field** of the corresponding database. |
| Step 2 | (PAN-DB only) Download the initial seed database and activate PAN-DB URL filtering. | 1. Click **Download** next to **Download Status** in the PAN-DB URL Filtering section.<br>2. Choose a region (North America, Europe, APAC, Japan) and then click **OK** to start the download.<br>3. After the download completes, click **Activate**.<br> <br>If PAN-DB is already the active URL filtering vendor and you click **Re-Download**, this will reactivate PAN-DB by clearing the dataplane and management plane caches and replacing them with the contents of the new seed database. You should avoid doing this unless it is necessary, as you will lose your cache, which is customized based on the web traffic that has previously passed through the firewall based on user activity. |
| Step 3 | (BrightCloud only) Enable cloud lookups for dynamically categorizing a URL if the category is not available on the local database. | 1. Access the CLI on the firewall.<br>2. Enter the following commands to enable Dynamic URL Filtering:<br>**configure**<br>**set deviceconfig setting url dynamic-url yes**<br>**commit** |

| Enable URL Filtering (Continued) | | |
|---|---|---|
| Step 4 | Configure Dynamic Updates for Applications and Threats and if you are using BrightCloud, configure the URL filtering updates.<br><br>The Applications and Threats updates may contain updates for URL filtering related to the Safe Search enforcement option available in the URL Filtering profile. For example, if Palo Alto Networks adds support for a new search provider vendor or if the method used to detect the Safe Search setting for an existing vendor changes, the update will be included in the Application and Threats updates.<br><br>BrightCloud updates include a database of approximately 20 million websites that are stored on the firewall drive, so the URL Filtering update must be scheduled to receive these updates.<br><br> A Threat Prevention license is required to receive content updates, which covers Antivirus and Applications and Threats. | 1. Select **Device > Dynamic Updates**.<br>2. In the **Applications and Threats** section, configure a schedule to receive updates periodically.<br>3. (BrightCloud only) In the **URL Filtering** section, configure a schedule to receive updates periodically. |

 © Palo Alto Networks, Inc.

# Determine URL Filtering Policy Requirements

The recommended practice for deploying URL filtering in your organization is to first start with a passive URL filtering profile that will alert on most categories. After setting the alert action, you can then monitor user web activity for a few days to determine the websites that are being accessed. After doing so, you can then make decisions on the websites and website categories that should be controlled.

- Configure and Apply a Passive URL Filtering Profile
- Monitor Web Activity

## Configure and Apply a Passive URL Filtering Profile

Because the default URL filtering profile blocks risky and threat-prone content, as a best practice, clone this profile when creating a new profile in order to preserve these default settings.

In the procedure that follows, threat-prone sites will be set to block and the other categories will be set to alert, which will cause all websites traffic to be logged. This may potentially create a large amount of log files, so it is best to do this for initial monitoring purposes to determine the types of websites your users are accessing. After determining the categories that your company approves of, those categories should then be set to allow, which will not generate logs. You can also reduce URL filtering logs by enabling the *Log container page only* option in the URL profile, so only the main page that matches the category will be logged, not subsequent pages/categories that may be loaded within the container page.

| **Configure and Apply a Passive URL Filtering Profile** | | |
|---|---|---|
| Step 1 | Clone the default URL filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select the default profile and then click **Clone**. The new profile will be named *default-1*.<br>3. Select the new profile and rename it. For example, rename it to *URL-Monitoring*. |

| Configure and Apply a Passive URL Filtering Profile (Continued) | | |
|---|---|---|
| Step 2 | Configure the action for all categories to *alert*, except for threat-prone categories, which should remain blocked.<br><br>Tip: To select all items in the category list from a Windows system, click the first category then hold down the shift key and click the last category—this will select all categories. Hold the control key (ctrl) down and click items that should be deselected. On a Mac, do the same using the shift and command keys. You could also just set all categories to alert and manually change the recommended categories back to block. | 1. In the section that lists all URL categories, select all categories and then deselect the following categories:<br>• abused-drugs<br>• adult<br>• gambling<br>• hacking<br>• malware<br>• phishing<br>• questionable<br>• weapons<br>2. To the right of the *Action* column heading, mouse over and select the down arrow and then select **Set Selected Actions** and choose **alert**.<br> 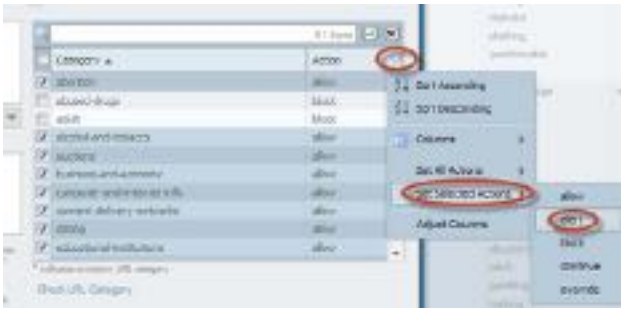 <br>3. Click **OK** to save. |
| Step 3 | Apply the URL profile to the security policy that allows web traffic for users. | 1. Select **Policies > Security** and select the appropriate security policy to modify it.<br>2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the new profile.<br>3. Click **OK** to save. |
| Step 4 | Save the configuration. | Click **Commit**. |
| Step 5 | View the URL filtering logs to determine all of the website categories that your users are accessing. In this example, some categories are set to block, so those categories will also appear in the logs.<br><br>For information on viewing the logs and generating reports, see Monitor Web Activity. | Select **Monitor > Logs > URL Filtering**. A log entry will be created for any website that exists in the URL filtering database that is in a category that is set to any action other than *allow*. |


© Palo Alto Networks, Inc.

## Monitor Web Activity

URL filtering logs and reports show all user web activity for URL categories that are set to alert, block, continue, or override. By monitoring the logs, you can gain a better understanding of the web activity of your user base to determine a web access policy.

This section assumes that a URL profile is configured as described in Configure and Apply a Passive URL Filtering Profile.

The following topics describe how to monitor web activity:

- Interpret the URL Filtering Logs
- Use the ACC to Monitor Web Activity
- View URL Filtering Reports
- Configure Custom URL Filtering Reports

### Interpret the URL Filtering Logs

The following bullet points show examples of the URL filtering logs (**Monitor > Logs > URL filtering**).

- **Alert log—**In this log the category is shopping and the action is alert.

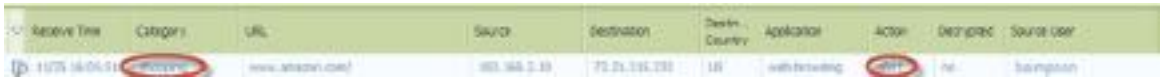

- **Block log—**In this log, the category alcohol-and-tobacco was set to block, so the action is block-url and the user will see a response page indicating that the website was blocked.

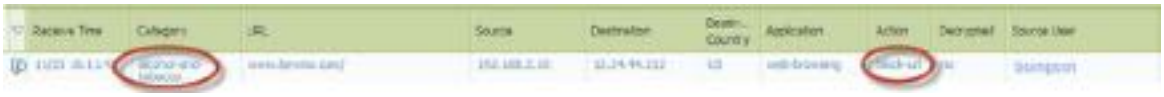

- **Alert log on encrypted website—**In this example, the category is social-networking and the application is facebook-base, which is required to access the Facebook website and other Facebook applications. Because faceboook.com is always encrypted using SSL, the traffic was decrypted by the firewall, which allows the website to be recognized and controlled if needed.

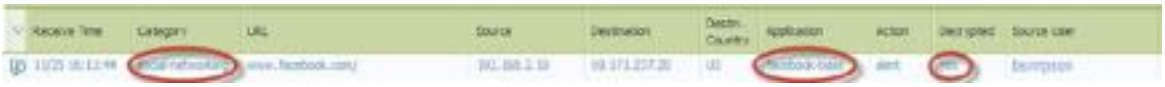

You can also add several other columns to your URL Filtering log view, such as: to and from zone, content type, and whether or not a packet capture was performed. To modify what columns to display, click the down arrow in any column and select the attribute to display.

To view the complete log details and/or request a category change for the given URL that was accessed, click the log details icon in the first column of the log.

## Use the ACC to Monitor Web Activity

For a quick view of the most common categories being accessed in your environment, select the **ACC** tab and scroll down to the **URL Filtering** section. Along the top of this window, you can also set the time range, sort by option, and define how many results will appear. Here you will see the most popular categories that are accessed by your users sorted by the most popular at the top of the list. In this example, *computer-and-internet-info* is the

 © Palo Alto Networks, Inc.

most accessed category, followed by *private-ip-addresses* (internal servers), and *search-engines*. In the drop-down in the upper right of the statistics, you can also choose to list by URL Categories, Blocked URL Categories, and Blocked URLs.

## View URL Filtering Reports

To view the default URL filtering reports, select **Monitor > Reports** and under the **URL Filtering Reports** section, choose one of the reports. You can generate reports on URL Categories, URL users, Web Sites accessed, Blocked Categories, and more. The reports are based on a 24 hour period and the day is selected by choosing a day in the calendar section. You can also export the report to PDF, CSV, or XML.

## View the User Activity Report

This report provides a quick method of viewing user or group activity and also provides an option to view browse time activity.

| Generate a User Activity Report | |
|---|---|
| Step 1 Configure a User Activity Report | 1. Select **Monitor > PDF Reports > User Activity Report**.<br>2. Enter a report **Name** and select the report type. Select **User** to generate a report for one person, or select **Group** for a group of users.<br>You mustEnable User-ID in order to be able to select user or group names. If User-ID is not configured, you can select the type **User** and enter the IP address of the user's computer.<br>3. Enter the Username/IP address for a user report or enter the group name for a user group report.<br>4. Select the time period. You can select an existing time period, or select **Custom**.<br>5. Select the **Include Detailed Browsing** check box, so browsing information is included in the report.<br> |
| Step 2 Run the user activity report and then download the report. | 1. Click **Run Now**.<br>2. After the report is generated, click the **Download User Activity Report** link.<br> 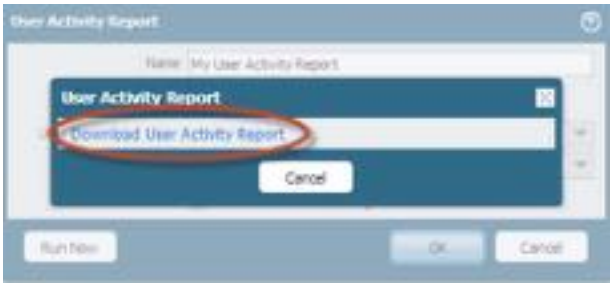 <br>3. After the report is downloaded, click **Cancel** and then click **OK** to save the report. |


© Palo Alto Networks, Inc.

**Generate a User Activity Report (Continued)**

Step 3 View the user activity report by opening the PDF file that was downloaded. The top of the report will contain a table of contents similar to the following:

User Activity Report for 192.168.2.10
Tuesday, December 31, 2013 09:35:47 - 10:35:46


Application Usage 2
Traffic Summary by URL Category 3
Browsing Summary by URL Category 4
Browsing Summary by Website 5
Blocked Browsing Summary by Website 6
Detailed Web Browsing Activity 7


Step 4 Click an item in the table of contents to view details. For example, click *Traffic Summary by URL Category* to view statistics for the selected user or group.

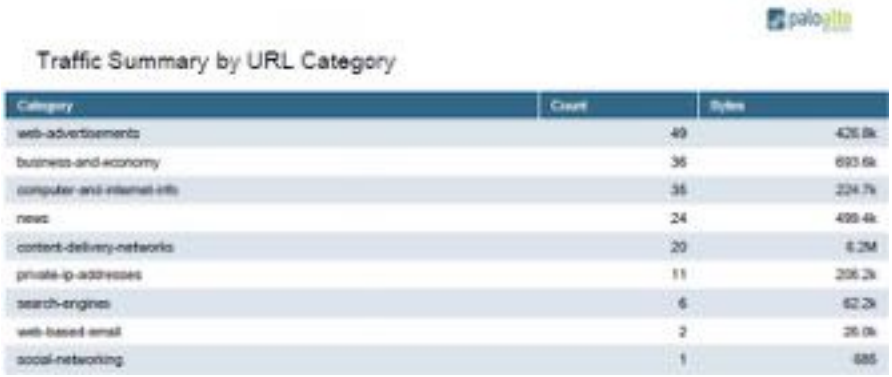
Traffic Summary by URL Category

| Category | Count | Bytes |
|---|---|---|
| web-advertisements | 49 | 426.8k |
| business-and-economy | 36 | 693.6k |
| computer-and-internet-info | 36 | 224.7k |
| news | 24 | 499.4k |
| content-delivery-networks | 20 | 6.2M |
| private-ip-addresses | 11 | 206.2k |
| search-engines | 6 | 62.2k |
| web-based-email | 2 | 26.0k |
| social-networking | 1 | 686 |

## Configure Custom URL Filtering Reports

To generate a detailed report that can also be scheduled, you can configure a custom report and select from a list of all available URL filtering log fields.

**Configure a Custom URL Filtering Report**

| | |
|---|---|
| Step 1 Add a new custom report. | 1. Select **Monitor > Manage Custom Reports** and click **Add**.<br>2. Enter a report name in the **Name** field. For example, *My-URL-Custom-Report*.<br>3. From the **Database** drop-down, select **URL Log**. |

| Configure a Custom URL Filtering Report (Continued) | | |
|---|---|---|
| Step 2 | Configure report options. | 1. Select the **Time Frame** drop-down and select a range.<br>2. (Optional) To customize how the report is sorted and grouped, select **Sort By** and chose the number of items to display (top 25 for example) and then select **Group By** and select an option such as **Category**, and then select how many groups will be defined.<br>3. In the **Available Columns** list, select the fields to include the report. The following columns are typically used for a URL report:<br>• Action<br>• Category<br>• Destination Country<br>• Source User<br>• URL<br> |
| Step 3 | Run the report to check the results. If the results are satisfactory, set a schedule to run the report automatically. | 1. Click the **Run Now** icon to immediately generate the report that will appear in a new tab.<br>2. (Optional) Click the **Schedule** check box to run the report once per day. This will generate a daily report that details web activity over the last 24 hours. To access the report, select **Monitor > Report** and then expand **Custom Reports** on the right column and select the report. |
| Step 4 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

## Define Website Controls

After you Determine URL Filtering Policy Requirements, you should have a basic understanding on what types of websites and website categories that your users are accessing. With this information, you are now ready to customize your URL filtering policies to control how your users access the web. The procedures that follow describe how to change the actions in URL profiles, use the Safe Search Enforcement option, and other features related to controlling content.

| Configure Website Controls | | |
|---|---|---|
| Step 1 | Customize the URL profile to control websites and website categories. | 1. Select **Objects > Security Profiles > URL Filtering** and modify your URL profile.<br>2. In the Category list, select the appropriate action for each URL category you want to control. For example, you may want to block categories such as auctions, gaming, and dating, but allow social-networking. |
| Step 2 | Configure websites that should always be blocked or allowed.<br><br>For example, to reduce URL filtering logs, you may want to enter all of your corporate websites in the allow list, so no log will be generated for those sites. If there is a website this is being overly used and is not work related in any way, you can add those in the block list.<br><br>Items in the block list will always be blocked regardless of the action for the given category, and URLs in the allow list will always be allowed.<br><br>For more information on the proper format and wildcards usage, see Block and Allow Lists. | 1. In the URL filtering profile, enter URLs or IP addresses in the **Block List** and select an action:<br>• **Block**—Block the URL.<br>• **Continue**—Users will be prompted with a response page when visiting a site that matches the defined URL category. If the user clicks *Continue*, the web page will open.<br>• **Override**—The user will be a prompted for a password to continue to the website.<br>• **Alert**—Allow the user to access the website and add an alert log entry in the URL log.<br>2. For the **Allow** list, enter IP addresses or URLs that should always be allowed. Each row must be separated by a new line.<br><br>In the following URL profile, social-networking is blocked, but Facebook is allowed. The corporate website paloaltonetworks.com is also allowed, so no logging will occur. The block list contains a server IP address that will always be blocked.<br><br> <br><br>3. (Optional) Enable Safe Search Enforcement. |
| Step 3 | Enable logging for Container Pages only. | Select the **Log container page only** check box. |

| Configure Website Controls | | |
|---|---|---|
| Step 4 | Save the URL filtering profile. | Click **OK**. |
| Step 5 | (Optional) Enable response pages on the ingress interface (the interface that first receives traffic for your users).<br><br>This option is required if you enable the *continue* action is configured for any URL category. | 1. Select **Network > Network Profiles > Interface Mgmt** and either add a new profile or edit an existing profile.<br>2. Click the **Response Pages** check box to enable.<br>3. Click **OK** to save the profile.<br>4. Select **Network > Interfaces** and then edit the layer 3 interface or VLAN interface that is your ingress interface.<br>5. Click the **Advanced** tab and select the Interface Mgmt profile that has the response page option enabled and select it from the drop-down menu.<br>6. Click **OK** to save the interface configuration. |
| Step 6 | (Optional) Customize URL filtering response pages.<br><br>There are three different response pages for URL filtering:<br>• **URL Filtering and Category Match Block Page**—Access blocked by a URL filtering profile or because the URL category is blocked by a security policy.<br>• **URL Filtering Continue and Override Page**—Page with initial block policy that allows users to bypass the block. With the override page, a password is required for the user to override the policy that blocks the URL.<br>• **URL Filtering Safe Search Block Page**—Access blocked by a security policy with a URL filtering profile that has the Safe Search Enforcement option enabled. The user will see this page if a search is performed using Google, Bing, or Yahoo and their browser or search engine account setting for Safe Search is not set to strict. | 1. Select **Device > Response Pages**.<br>2. Click the URL filtering response page that you would like to modify.<br>3. Select the response page (predefined or shared) and then click the **Export** link and save the file to your desktop.<br> Predefined is the default response page and shared is a custom response page created by an administrator.<br>4. Modify the response page using a text editor and then save the file.<br>5. From the response page dialogue, click **Import** and select the newly modified response page and then click **OK** to import the file.<br>6. The newly imported response page will become the active response page. |

 © Palo Alto Networks, Inc.

| Configure Website Controls | | |
|---|---|---|
| Step 7 | Configure a URL filtering override to create a temporary password that can be used by specific users to access sites that are blocked. | 1. Select **Device > Setup > Content ID**.<br>2. In the **URL Admin Override** section, click **Add** to configure a password.<br>3. (Optional) Set a custom override period by entering a new value in the **URL Admin Override Timeout** field. By default, users can access a blocked URL categories for 15 minutes.<br>4. When setting the password, you can choose **Transparent** or **Redirect**.<br>• **Transparent**—The firewall intercepts the browser traffic and impersonates the original destination URL, issuing an HTTP 401 to invoke authentication. However, because the firewall does not have the real certificate for the destination URL, the browser will display a certificate error to users attempting to access a secure site. Therefore you should only use this mode when absolutely necessary, such as in Layer 2 or virtual wire deployments.<br>• **Redirect**—The firewall intercepts unknown HTTP or HTTPS sessions and redirects them to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to perform authentication. This is the preferred mode because it provides a better end-user experience (no certificate errors). However, it does require additional Layer 3 configuration. Another benefit of the Redirect mode is that it provides the use of session cookies, which enable the user to continue browsing to authenticated sites without requiring re-mapping each time the time-out expires. |
| Step 8 | Save the configuration.<br> To test the URL filtering configuration, simply access a website in a category that is set to block or continue to see if the appropriate action is performed. | Click **Commit**. |

# Enable Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos for search query return traffic. You can configure Safe Search Enforcement on the Palo Alto Networks next-generation firewall to prevent search requests that do not have the strictest safe search settings enabled.

The Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address is not compatible with Safe Search Enforcement on the firewall.

There are two ways to enforce Safe Search on the firewall:

- Block Search Results that are not Using Strict Safe Search Settings
- Enable Transparent Safe Search Enforcement

## Block Search Results that are not Using Strict Safe Search Settings

By default, when you enable safe search enforcement, when a user attempts to perform a search without using the strictest safe search settings, the firewall will block the search query results and display the URL Filtering Safe Search Block Page. This page provides a link to the search settings page for the corresponding search provider so that the end user can enable the safe search settings. If you plan to use this default method for enforcing safe search, you should communicate the policy to your end users prior to deploying the policy. See Table: Search Provider Safe Search Settings for details on how each search provider implements safe search. The default URL Filtering Safe Search Block Page provides a link to the search settings for the corresponding search provider. You can optionally customize the URL filtering response pages.

Alternatively, to enable safe search enforcement so that it is transparent to your end users, configure the firewall to Enable Transparent Safe Search Enforcement.

 © Palo Alto Networks, Inc.

| Enable Safe Search Enforcement | | |
|---|---|---|
| Step 1 | Enable Safe Search Enforcement in the URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select an existing profile to modify, or clone the default profile to create a new profile.<br>3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.<br>4. (Optional) Restrict users to specific search engines:<br>a. On the **Categories** tab, set the **search-engines** category to **block**.<br>b. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:<br>`www.google.com`<br>`www.bing.com`<br>5. Configure other settings as necessary to Define Website Controls.<br>6. Click **OK** to save the profile. |
| Step 2 | Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet. | 1. Select **Policies > Security** and select a rule to which to apply the URL filtering profile that you just enabled for Safe Search Enforcement.<br>2. On the **Actions** tab, select the **URL Filtering** profile.<br>3. Click **OK** to save the security policy rule. |
| Step 3 | Enable SSL Forward Proxy decryption.<br>Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings. | 1. Add a custom URL category for the search sites:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as SearchEngineDecryption.<br>c. **Add** the following to the Sites list:<br>`www.bing.*`<br>`www.google.*`<br>`search.yahoo.*`<br>d. Click **OK** to save the custom URL category object.<br>2. Follow the steps to Configure SSL Forward Proxy.<br>3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**. |

| Enable Safe Search Enforcement (Continued) | | |
|---|---|---|
| Step 4 | (Optional, but recommended) Block Bing search traffic running over SSL.<br><br>Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL. | 1. Add a custom URL category for Bing:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as EnableBingSafeSearch.<br>c. **Add** the following to the Sites list:<br>`www.bing.com/images/*`<br>`www.bing.com/videos/*`<br>d. Click **OK** to save the custom URL category object.<br>2. Create another URL filtering profile to block the custom category you just created:<br>a. Select **Objects > Security Profiles > URL Filtering**.<br>b. **Add** a new profile and give it a descriptive **Name**.<br>c. Locate the custom category in the Category list and set it to **block**.<br>d. Click **OK** to save the URL filtering profile.<br>3. Add a security policy rule to block Bing SSL traffic:<br>a. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.<br>b. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.<br>c. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.<br>d. Select **TCP** as the **Protocol** and set the **Destination Port** to `443`.<br>e. Click **OK** to save the rule.<br>f. Use the **Move** options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled. |
| Step 5 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

| Enable Safe Search Enforcement (Continued) | |
|---|---|
| Step 6 Verify the Safe Search Enforcement configuration.<br><br>This verification step only works if you are using block pages to enforce safe search. If you are using transparent safe search enforcement, the firewall block page will invoke a URL rewrite with the safe search parameters in the query string. | 1. From a computer that is behind the firewall, disable the strict search settings for one of the supported search providers. For example, on bing.com, click the **Preferences** icon on the Bing menu bar.<br><br> 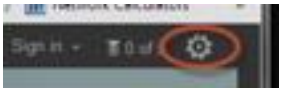 <br><br>2. Set the **SafeSearch** option to **Moderate** or **Off** and click **Save**.<br>3. Perform a Bing search and verify that the URL Filtering Safe Search Block page displays instead of the search results:<br><br> <br><br>4. Use the link in the block page to go to the search settings for the search provider and set the safe search setting back to the strictest setting (**Strict** in the case of Bing) and then click **Save**.<br>5. Perform a search again from Bing and verify that the filtered search results display instead of the block page. |

## Enable Transparent Safe Search Enforcement

If you want to enforce filtering of search query results with the strictest safe search filters, but you don't want your end users to have to manually configure the settings, you can enable transparent safe search enforcement as follows. This functionality is supported on Google, Yahoo, and Bing search engines only and requires Content Release version 475 or later.

| **Enable Transparent Safe Search Enforcement** | | |
|---|---|---|
| Step 1 | Make sure the firewall is running Content Release version 475 or later. | 1. Select **Device > Dynamic Updates**.<br>2. Check the **Applications and Threats** section to determine what update is currently running.<br>3. If the firewall is not running the required update or later, click **Check Now** to retrieve a list of available updates.<br>4. Locate the required update and click **Download**.<br>5. After the download completes, click **Install**. |
| Step 1 | Enable Safe Search Enforcement in the URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select an existing profile to modify, or clone the default profile to create a new one.<br>3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.<br>4. (Optional) Allow access to specific search engines only:<br>a. On the **Categories** tab, set the **search-engines** category to **block**.<br>b. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:<br>`www.google.com`<br>`www.bing.com`<br>5. Configure other settings as necessary to Define Website Controls.<br>6. Click **OK** to save the profile. |
| Step 2 | Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet. | 1. Select **Policies > Security** and select a rule to which to apply the URL filtering profile that you just enabled for Safe Search Enforcement.<br>2. On the **Actions** tab, select the **URL Filtering** profile.<br>3. Click **OK** to save the security policy rule. |


© Palo Alto Networks, Inc.

| Enable Transparent Safe Search Enforcement (Continued) | |
|---|---|
| Step 3 (Optional, but recommended) Block Bing search traffic running over SSL.<br><br>Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL. | 1. Add a custom URL category for Bing:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as EnableBingSafeSearch.<br>c. **Add** the following to the Sites list:<br>`www.bing.com/images/*`<br>`www.bing.com/videos/*`<br>d. Click **OK** to save the custom URL category object.<br>2. Create another URL filtering profile to block the custom category you just created:<br>a. Select **Objects > Security Profiles > URL Filtering**.<br>b. **Add** a new profile and give it a descriptive **Name**.<br>c. Locate the custom category you just created in the Category list and set it to **block**.<br>d. Click **OK** to save the URL filtering profile.<br>3. **Add** a security policy rule to block Bing SSL traffic:<br>a. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.<br>b. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.<br>c. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.<br>d. Select **TCP** as the **Protocol**, set the **Destination Port** to `443`.<br>e. Click **OK** to save the rule.<br>f. Use the **Move** options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled. |

© Palo Alto Networks, Inc.

**Enable Transparent Safe Search Enforcement (Continued)**

**Step 4** Edit the URL Filtering Safe Search Block Page, replacing the existing code with the Javascript for rewriting search query URLs to enforce safe search transparently.

1. Select **Device > Response Pages > URL Filtering Safe Search Block Page**.
2. Select **Predefined** and then click **Export** to save the file locally.
3. Use an HTML editor and replace all of the existing block page text with the following text and then save the file:

```
<html>
<head>
<script>

var s_u = location.href;

//bing
b_n = s_u.search("www.bing.com/")

if (b_n > 0) {
    s_u = s_u + "&adlt=strict";
}

//google
g_n = s_u.search("www.google.com/")

if (g_n > 0) {
    s_u = s_u + "&safe=active";
}

// yahoo
y_n = s_u.search("search.yahoo.com")

if (y_n > 0) {
    s_u = s_u.replace(/&vm=p/ig,"");
    s_u = s_u + "&vm=r";
}

window.location.replace(s_u);

document.getElementById("java_off").innerHTML = 'You are
  being redirected to a safer search!';

</script>
<title>Search Blocked</title>
<meta http-equiv="Content-Type" content="text/html;
  charset=utf-8">
<META HTTP-EQUIV="PRAGMA" CONTENT="NO-CACHE">
<meta name="viewport" content="initial-scale=1.0">
</head>
    <body bgcolor="#e7e8e9" id="java_off">
    Your search is not safesearch enforced!
    Please enable Java script in your browser.
</body>
</html>
```

**Step 5** Import the edited URL Filtering Safe Search Block page onto the firewall.

1. To import the edited block page, select **Device > Response Pages > URL Filtering Safe Search Block Page**.
2. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.
3. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select **shared** to make it available to all virtual systems.
4. Click **OK** to import the file.


© Palo Alto Networks, Inc.

| Enable Transparent Safe Search Enforcement (Continued) | | |
|---|---|---|
| Step 6 | Enable SSL Forward Proxy decryption.<br><br>Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings. | 1. Add a custom URL category for the search sites:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as SearchEngineDecryption.<br>c. **Add** the following to the Sites list:<br>`www.bing.*`<br>`www.google.*`<br>`search.yahoo.*`<br>d. Click **OK** to save the custom URL category object.<br>2. Follow the steps to Configure SSL Forward Proxy.<br>3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**. |
| Step 7 | Save the configuration. | Click **Commit**. |

# URL Filtering Use Case Examples

The following use cases show how to use App-ID to control a specific set of web-based applications and how to use URL categories as a match criteria in a policy. When working with App-ID, it is important to understand that each App-ID signature may have dependencies that are required to fully control an application. For example, with Facebook applications, the App-ID facebook-base is required to access the Facebook website and to control other Facebook applications. For example, to configure the firewall to control Facebook email, you would have to allow the App-IDs facebook-base and facebook-mail. As another example, if you search Applipedia (the App-ID database) for LinkedIn, you will see that in order to control LinkedIn mail, you need to apply the same action to both App-IDs: linkedin-base and linkedin-mail. To determine application dependencies for App-ID signatures, visit Applipedia, search for the given application, and then click the application for details.

The User-ID feature is required to implement policies based on users and groups and a Decryption policy is required to identify and control websites that are encrypted using SSL.

This section includes two uses cases:

- Use Case: Control Web Access
- Use Case: Use URL Categories in Policy

 © Palo Alto Networks, Inc.

## Use Case: Control Web Access

When using URL filtering to control user website access, there may be instances where granular control is required for a given website. In this use case, a URL filtering policy is applied to the security policy that allows web access for your users and the *social-networking* URL category is set to block, but the allow list in the URL profile is configured to allow the social networking site Facebook. To further control Facebook, the company policy also states that only marketing has full access to Facebook and all other users within the company can only read Facebook posts and cannot use any other Facebook applications, such as email, posting, chat, and file sharing. To accomplish this requirement, App-ID must be used to provide granular control over Facebook.

The first security rule will allow marketing to access the Facebook website as well as all Facebook applications. Because this allow rule will also allow access to the Internet, threat prevention profiles are applied to the rule, so traffic that matches the policy will be scanned for threats. This is important because the allow rule is terminal and will not continue to check other rules if there is a traffic match.

| **Control Web Access** | | |
|---|---|---|
| Step 1 | Confirm that URL filtering is licensed. | 1. Select **Device > Licenses** and confirm that a valid date appears for the URL filtering database that will used. This will either be PAN-DB or BrightCloud.<br>2. If a valid license is not installed, see Enable URL Filtering. |
| Step 2 | Confirm that User-ID is working. User-ID is required to create policies based on users and groups. | 1. To check group mapping, from the CLI, enter the following command:<br>`show user group-mapping statistics`<br>2. To check user mapping, from the CLI, enter the following command:<br>`show user ip-user-mapping-mp all`<br>3. If statistics do not appear and/or IP to user mapping information is not displayed, see User-ID. |
| Step 3 | Set up a URL filtering profile by cloning the default profile. | 1. Select **Objects > Security Profiles > URL Filtering** and select the *default* profile.<br>2. Click the **Clone** icon. A new profile should appear named *default-1*.<br>3. Select the new profile and rename it. |

| Control Web Access (Continued) | | |
|---|---|---|
| Step 4 | Configure the URL filtering profile to block *social-networking* and allow *Facebook*. | 1. Modify the new URL filtering profile and in the **Category** list scroll to *social-networking* and in the **Action** column click on *allow* and change the action to *block*.<br>2. In the **Allow List** box, type *facebook.com*, press enter to start a new line and then type **.facebook.com*. Both of these formats are required, so all URL variants a user may use will be identified, such as facebook.com, www.facebook.com, and https://facebook.com.<br> 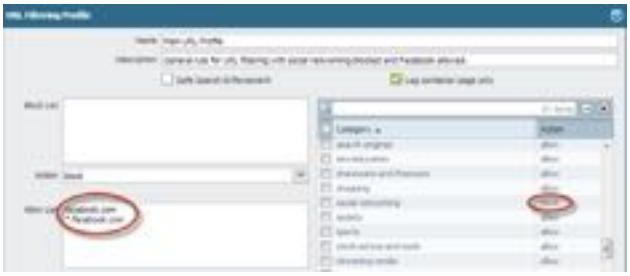 <br>3. Click **OK** to save the profile. |
| Step 5 | Apply the new URL filtering profile to the security policy rule that allows web access from the user network to the Internet. | 1. Select **Policies > Security** and click on the policy rule that allows web access.<br>2. On the **Actions** tab, select the URL profile you just created from the **URL Filtering** drop down.<br> 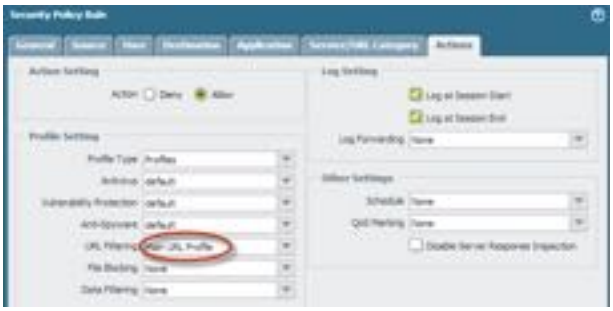 <br>3. Click **OK** to save. |


© Palo Alto Networks, Inc.

**Control Web Access (Continued)**

Step 6 Create the security policy that will allow marketing access the Facebook website and all Facebook applications.

This rule must precede other rules because it is more specific than the other policies and because it is an allow rule, which will terminate when a traffic match occurs.

1. Select **Policies > Security** and click **Add**.
2. Enter a **Name** and optionally a **Description** and **Tag**(s).
3. On the **Source** tab add the zone where the users are connected.
4. On the **User** tab in the **Source User** section click **Add**.
5. Select the directory group that contains your *marketing* users.
6. On the **Destination** tab, select the zone that is connected to the Internet.
7. On the **Applications** tab, click **Add** and add the *facebook* App-ID signature.
8. On the **Actions** tab, add the default profiles for **Antivirus**, **Vulnerability Protection**, and **Anti-Spyware**.

9. Click **OK** to save the security profile.

The *facebook* App-ID signature used in this policy encompasses all Facebook applications, such as facebook-base, facebook-chat, and facebook-mail, so this is the only App-ID signature required in this rule.

With this policy in place, when a marketing employee attempts to access the Facebook website or any Facebook application, the rule matches based on the user being part of the marketing group. For traffic from any user outside of marketing, the rule will be skipped because there would not be a traffic match and rule processing would continue.

**Control Web Access (Continued)**

Step 7 Configure the security policy to block all other users from using any Facebook applications, other than simple web browsing. The easiest way to do this is to clone the marketing allow policy and then modify it.

1. From **Policies > Security** click the marketing Facebook allow policy you created earlier to highlight it and then click the **Clone** icon.
2. Enter a **Name** and optionally enter a **Description** and **Tag**('s).
3. On the **User** tab highlight the marketing group and delete it and in the drop-down select **any**.
4. On the **Applications** tab, click the *facebook* App-ID signature and delete it.
5. Click **Add** and add the following App-ID signatures:
   - facebook-apps
   - facebook-chat
   - facebook-file-sharing
   - facebook-mail
   - facebook-posting
   - facebook-social-plugin
6. On the **Actions** tab in the **Action Setting** section, select **Deny**. The profile settings should already be correct because this rule was cloned.

7. Click **OK** to save the security profile.
8. Ensure that this new deny rule is listed after the marketing allow rule, to ensure that rule processing occurs in the correct order to allow marketing users and then to deny/limit all other users.
9. Click **Commit** to save the configuration.

With these policies in place, any user who is part of the marketing group will have full access to all Facebook applications and any user that is not part of the marketing group will only have read-only access to the Facebook website and will not be able to use Facebook functions such as post, chat, email, and file sharing.

 © Palo Alto Networks, Inc.

## Use Case: Use URL Categories in Policy

URL categories can also be used as a match criteria in the following policy types: captive portal, decryption, security, and QoS. In this use case, URL categories will be used in decryption policies to control which web categories should be decrypted or not decrypted. The first rule is a no-decrypt rule that will not decrypt user traffic if the website category is *financial-services* or *health-and-medicine* and the second rule will decrypt all other traffic. The decryption policy type is *ssl-forward-proxy*, which is used for controlling decryption for all outbound connections performed by users.

**Configure a Decryption Policy Based on URL Category**

Step 1 Create the no-decrypt rule that will be listed first in the decryption policies list. This will prevent any website that is in the *financial-services* or *health-and-medicine* URL categories from being decrypted.

1. Select **Policies > Decryption** and click **Add**.
2. Enter a **Name** and optionally enter a **Description** and **Tag**(s).
3. On the **Source** tab add the zone where the users are connected.
4. On the **Destination** tab enter the zone that is connected to the Internet.
5. On the **URL Category** tab click **Add** and select the *financial-services* and *health-and-medicine* URL categories.
6. On the **Options** tab, set the action to **No Decrypt** and the **Type** to **SSL Forward Proxy**.

7. Click **OK** to save the policy.

Step 2 Create the decryption policy that will decrypt all other traffic. This policy will be listed after the no-decrypt policy.

1. Select the no-decrypt policy you created previously and then click **Clone**.
2. Enter a **Name** and optionally enter a **Description** and **Tag**(s).
3. On the **URL Category** tab, select *financial-services* and *health-and-medicine* and then click the **Delete** icon.
4. On the **Options** tab, set the action to **Decrypt** and the **Type** to **SSL Forward Proxy**.

5. Ensure that this new decryption rule is listed after the no-decrypt rule as shown in the previous screen capture. This will ensure that rule processing occurs in the correct order, so websites in the *financial-services* and *health-and-medicine* are not decrypted
6. Click **OK** to save the policy.

| Configure a Decryption Policy Based on URL Category (Continued) | | |
|---|---|---|
| Step 3 | (BrightCloud only) Enable cloud lookups for dynamically categorizing a URL when it the category is not available on the local database on the firewall. | 1. Access the CLI on the firewall.<br>2. Enter the following commands to enable Dynamic URL Filtering:<br>a. **configure**<br>b. **set deviceconfig setting url dynamic-url yes**<br>c. commit |
| Step 4 | Save the configuration. | Click **Commit**. |

With these two decrypt policies in place, any traffic destined for the *financial-services* or *health-and-medicine* URL categories, the traffic will not be decrypted. All other traffic will be decrypted.

You can also define more granular control over decryption policies by defining decryption profiles, which are used to perform checks such as server certificate checks or blocking sessions with expired certificates. The profile is then added in the **Options** tab of the decryption policy. For a complete list of checks that can be performed, select **Objects > Decryption Profiles** from the firewall and then click the help icon.

Now that you have a basic understanding of the powerful features of URL filtering, App-ID, and User-ID, you can apply similar policies to your firewall to control any application in the Palo Alto Networks App-ID signature database and control any website contained in the URL filtering database.

For help in troubleshooting URL filtering issues, see Troubleshoot URL Filtering.

 © Palo Alto Networks, Inc.

# Troubleshoot URL Filtering

This following topics provide troubleshooting guidelines for diagnosing and resolving common URL filtering problems.

- Problems Activating PAN-DB
- PAN-DB Cloud Connectivity Issues
- URLs Classified as Not-Resolved
- Incorrect Categorization
- URL Database Out of Date

## Problems Activating PAN-DB

The following table describes procedures that you can use to resolve issues with activating PAN-DB.

| Troubleshoot PAN-DB Activation Issues | |
|---|---|
| Step 1 | Access the CLI on the firewall. |
| Step 2 | Verify if PAN-DB has been activated by running the following command:<br>`admin@PA-200> show system setting url-database`<br>If the response is `paloaltonetworks`, then PAN-DB is the active vendor. |
| Step 3 | Verify that the firewall has a valid PAN-DB license by running the following command:<br>`admin@PA-200> request license info`<br>You should see the license entry Feature: PAN_DB URL Filtering. If the license is not installed, you will need to obtain and install a license. See Configure URL Filtering. |
| Step 4 | After the license is installed, download a new PAN-DB seed database by running the following command:<br>`admin@PA-200> request url-filtering download paloaltonetworks region` *<region>* |
| Step 5 | Check the download status by running the following command:<br>`admin@PA-200> request url-filtering download status vendor paloaltonetworks`<br>• If the message is different than `PAN-DB download: Finished successfully`, stop here, there may be a problem connecting to the cloud. Attempt to solve the connectivity issue by performing basic network troubleshooting between the firewall and the Internet. For more information, see PAN-DB Cloud Connectivity Issues.<br>• If the message is `PAN-DB download: Finished successfully`, the firewall successfully downloaded the URL seed database. Try to enable PAN-DB again by running the following command:<br>`admin@PA-200> set system setting url-database paloaltonetworks` |
| Step 6 | If the problems persists, contact Palo Alto Networks support. |

## PAN-DB Cloud Connectivity Issues

To check cloud connectivity, run the following command:

```
admin@pa-200> show url-cloud status
```

If the cloud is accessible, the expected response should be similar to the following:

```
admin@PA-200> show url-cloud status
PAN-DB URL Filtering
License :                          valid
Current cloud server :             s0000.urlcloud.paloaltonetworks.com
Cloud connection :                 connected
URL database version - device :    2013.11.18.000
URL database version - cloud :     2013.11.18.000  ( last update time
2013/11/19
 13:20:51 )
URL database status :              good
URL protocol version - device :    pan/0.0.2
URL protocol version - cloud :     pan/0.0.2
Protocol compatibility status :    compatible
```

If the cloud is note accessible, the expected response will be similar to the following:

```
admin@PA-200> show url-cloud status
PAN-DB URL Filtering
License :                          valid
Cloud connection :                 not connected
URL database version - device :    2013.11.18.000
URL database version - cloud :     2013.11.18.000  ( last update time
2013/11/19
 13:20:51 )
URL database status :              good
URL protocol version - device :    pan/0.0.2
URL protocol version - cloud :     pan/0.0.2
Protocol compatibility status :    compatible
```

The following table describes procedures that you can use to resolve issues based on the output of the show cloud status command, how to ping the URL cloud servers, and what to check if the firewall is in a High Availability (HA) configuration.

 © Palo Alto Networks, Inc.

| Troubleshoot Cloud Connectivity Issues |
|---|
| • PAN-DB URL Filtering license field shows invalid—Obtain and install a valid PAN-DB license.<br>• URL database status is out of date—Download a new seed database by running the following command:<br>• admin@pa-200> request url-filtering download paloaltonetworks region <region><br>• URL protocol version shows not compatible—Upgrade PAN-OS to the latest version.<br>• Attempt to ping the PAN-DB cloud server from the firewall by running the following command:<br>• admin@pa-200> ping source *ip-address* host s0000.urlcloud.paloaltonetworks.com<br>• For example, if your management interface IP address is 10.1.1.5, run the following command:<br>• admin@pa-200> ping source 10.1.1.5 host s0000.urlcloud.paloaltonetworks.com<br>• If the firewall is in an HA configuration, verify that the HA state of the devices supports connectivity to the cloud systems. You can determine the HA state by running the following command:<br>• admin@pa-200> show high-availability state<br>Connection to the cloud will be blocked if the firewall is not in one of the following states:<br>• active<br>• active-primary<br>• active-secondary<br>If the problem persists, contact Palo Alto Networks support. |

## URLs Classified as Not-Resolved

The following table describes procedures you can use to resolve issues where some or all of the URLs being identified by PAN-DB are classified as *Not-resolved*:

| Troubleshoot URLs Classified as Not-Resolved | |
|---|---|
| Step 1 | Check the PAN-DB cloud connection by running the following command:<br>`admin@PA-200> show url-cloud status`<br>The Cloud connection: field should show *connected*. If you see anything other than *connected*, any URL that do not exist in the management plane cache will be categorized as *not-resolved*. To resolve this issue, see PAN-DB Cloud Connectivity Issues. |
| Step 2 | If the cloud connection status shows *connected*, check the current utilization of the firewall. If the firewall's performance is spiking, URL requests may be dropped (may not reach the management plane), and will be categorized as *not-resolved*.<br>To view system resources, run the following command and view the *%CPU* and *%MEM* columns:<br>`admin@PA-200> show system resources`<br>You can also view system resources from the firewall's web interfaces by clicking the **Dashboard** tab and viewing the **System Resources** section. |
| Step 3 | If the problem persist, contact Palo Alto Networks support. |

## Incorrect Categorization

The following steps describe the procedures you can use if you identify a URL that does not have the correct categorization. For example, if the URL *paloaltonetworks.com* was categorized as *alcohol-and-tobacco*, the categorization is not correct; the category should be *computer-and-internet-info*.

| Troubleshoot Incorrect Categorization Issues | |
|---|---|
| Step 1 | Verify the category in the dataplane by running the following command:<br>`admin@PA-200>` show running url <*URL*><br>For example, to view the category for the Palo Alto Networks website, run the following command:<br>`admin@PA-200> show running url paloaltonetworks.com`<br>If the URL stored in the dataplane cache has the correct category (computer-and-internet-info in this example), then the categorization is correct and no further action is required. If the category is not correct, continue to the next step. |
| Step 2 | Verify if the category in the management plane by running the command:<br>`admin@PA-200> test url-info-host` <*URL*><br>For example:<br>`admin@PA-200> test url-info-host paloaltonetworks.com`<br>If the URL stored in the management plane cache has the correct category, remove the URL from the dataplane cache by running the following command:<br>`admin@PA-200> clear url-cache url` <*URL*><br>Next time the device requests the category for this URL, the request will be forwarded to the management plane. This will resolve the issue and no further action is required. If this does not solve the issue, go to the next step to check the URL category on the cloud systems. |
| Step 3 | Verify the category in the cloud by running the following command:<br>`admin@PA-200> test url-info-cloud <URL>` |
| Step 4 | If the URL stored in the cloud has the correct category, remove the URL from the dataplane and the management plane caches.<br>Run the following command to delete a URL from the dataplane cache:<br>`admin@PA-200> clear url-cache url <URL>`<br>Run the following command to delete a URL from the management plane cache:<br>`admin@PA-200> delete url-database url <URL>`<br>Next time the device queries for the category of the given URL, the request will be forwarded to the management plane and then to the cloud. This should resolve the category lookup issue. If problems persist, see the next step to submit a categorization change request. |
| Step 5 | To submit a change request from the web interface, go to the URL log and select the log entry for the URL you would like to have changed. |

 © Palo Alto Networks, Inc.

**Troubleshoot Incorrect Categorization Issues**

Step 6 Click the **Request Categorization** change link and follow instructions. You can also request a category change at the Palo Alto Networks Test A Site website by searching for the URL and then clicking the **Request Change** icon. To view a list of all available categories with descriptions of each category, refer to https://urlfiltering.paloaltonetworks.com/CategoryList.aspx.

If your change request is approved, you will receive an email notification. You then have two options to ensure that the URL category is updated on the firewall:

- Wait until the URL in the cache expires and the next time the URL is accessed by a user, the new categorization update will be put in the cache.
- Run the following command to force an update in the cache:

```
admin@PA-200> request url-filtering update url <URL>
```

## URL Database Out of Date

If you have observed through the syslog or the CLI that PAN-DB is out-of-date, it means that the connection from the firewall to the URL Cloud is blocked. This usually occurs when the URL database on the firewall is too old (version difference is more than three months) and the cloud cannot update the firewall automatically. In order to resolve this issue, you will need to re-download an initial seed database from the cloud (this operation is not blocked). This will result in an automatic re-activation of PAN-DB.

To manually update the database, perform one of the following steps:

– From the web interface, select **Device > Licenses** and in the **PAN-DB URL Filtering** section click the **Re-Download** link.

– From the CLI, run the following command:

```
admin@PA-200> request url-filtering download paloaltonetworks region <region_name>
```

When the seed database is re-download, the URL cache in the management plane and dataplane will be purged. The management plane cache will then be re-populated with the contents of the new seed database.

 © Palo Alto Networks, Inc.

# Quality of Service

Quality of Service (QoS) is a set of technologies that work on a network to guarantee its ability to dependably run high-priority applications and traffic under limited network capacity. QoS technologies accomplish this by providing differentiated handling and capacity allocation to specific flows in network traffic. This enables the network administrator to assign the order in which traffic is handled, and the amount of bandwidth afforded to traffic.

Palo Alto Networks Application Quality of Service (QoS) provides basic QoS applied to networks and extends it to provide QoS to applications and users.

Use the following topics to learn about and configure Palo Alto Networks Application QoS:

- QoS Overview
- QoS Concepts
- Configure QoS
- Configure QoS for a Virtual System
- QoS Use Case Examples

# QoS Overview

Use QoS to prioritize and adjust quality aspects of network traffic. You can assign the order in which packets are handled and allot bandwidth, ensuring preferred treatment and optimal levels of performance are afforded to selected traffic, applications, and users.

Service quality measurements subject to a QoS implementation are bandwidth (maximum rate of transfer), throughput (actual rate of transfer), latency (delay), and jitter (variance in latency). The capability to shape and control these service quality measurements makes QoS of particular importance to high-bandwidth, real-time traffic such as voice over IP (VoIP), video conferencing, and video-on-demand that has a high sensitivity to latency and jitter. Additionally, use QoS to achieve outcomes such as the following:

- Prioritize network and application traffic, guaranteeing high priority to important traffic or limiting non-essential traffic.
- Achieve equal bandwidth sharing among different subnets, classes, or users in a network.
- Allocate bandwidth externally or internally or both, applying QoS to both upload and download traffic or to only upload or download traffic.
- Ensure low latency for customer and revenue-generating traffic in an enterprise environment.
- Perform traffic profiling of applications to ensure bandwidth usage.

Each firewall model supports a maximum number of ports that can be configured with QoS.

QoS implementation on a Palo Alto Networks firewall begins with three primary configuration components that support a full QoS solution:

- QoS Profile
- QoS Policy
- QoS on the physical interface

Each of these options in the QoS configuration task facilitate a broader process that optimizes and prioritizes the traffic flow and allocates and ensures bandwidth according to configurable parameters.

Figure: QoS Traffic Flow shows traffic as it flows from the source, is shaped by the firewall with QoS enabled, and is ultimately prioritized and delivered to its destination.

**Figure: QoS Traffic Flow**

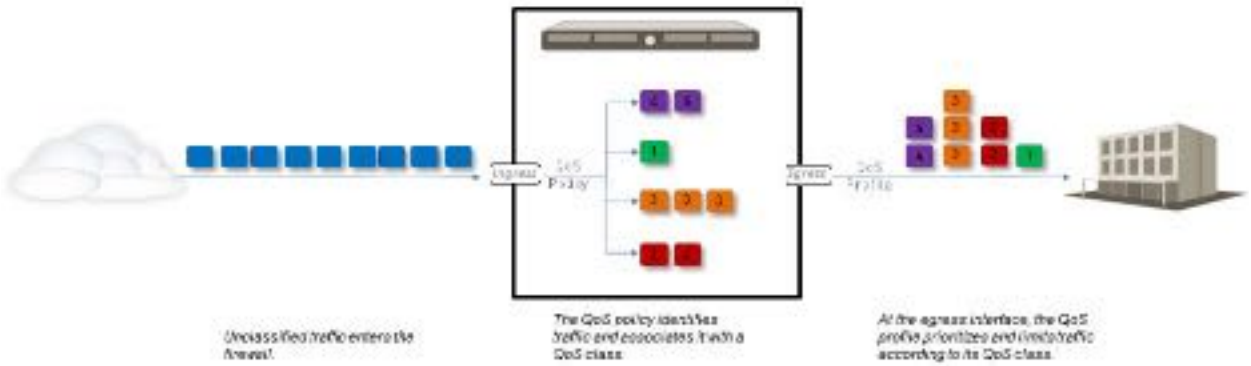

 © Palo Alto Networks, Inc.

The QoS configuration options allow you to control the traffic flow and define it at different points in the flow. The Figure: QoS Traffic Flow indicates where the configurable options define the traffic flow. Use the QoS Profile to define QoS classes and use the QoS Policy to associate QoS classes with selected traffic. Enable the QoS Profile on a physical interface to shape traffic according to the QoS configuration as it flows through the network.

You can configure a QoS Profile and QoS Policy individually or in any order, according to your preference. Each of the QoS configuration options has components that influence the definition of the other options and the QoS configuration options can be used to create a full and granular QoS policy or can be used sparingly with minimal administrator action.

© Palo Alto Networks, Inc.

# QoS Concepts

Use the following topics to learn about the different components and mechanisms of a QoS configuration on a Palo Alto Networks firewall:

- QoS for Applications and Users
- QoS Profile
- QoS Classes
- QoS Policy
- QoS Egress Interface
- QoS Cleartext and Tunneled Traffic

## QoS for Applications and Users

A Palo Alto Networks firewall provides basic QoS, controlling traffic leaving the firewall according to network or subnet, and extends the power of QoS to also classify and shape traffic according to application and user. The Palo Alto Networks firewall provides this capability by integrating the features App-ID and User-ID with the QoS configuration. App-ID and User-ID entries that exist to identify specific applications and users in your network are available in the QoS configuration so that you can easily specify applications and users to apply QoS to.

You can use a QoS Policy in the web interface (**Policies** > **QoS**) to apply QoS specifically to an application's traffic:

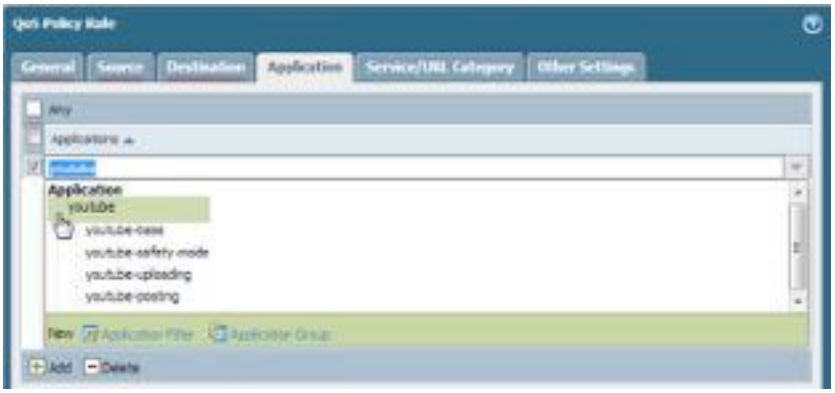


Or to a user's traffic:

See App-ID and User-ID for more information on these features.


© Palo Alto Networks, Inc.

## QoS Profile

Use a QoS profile to define values of up to eight QoS classes contained within that single profile (**Network** > **Network Profiles** > **QoS Profile**):

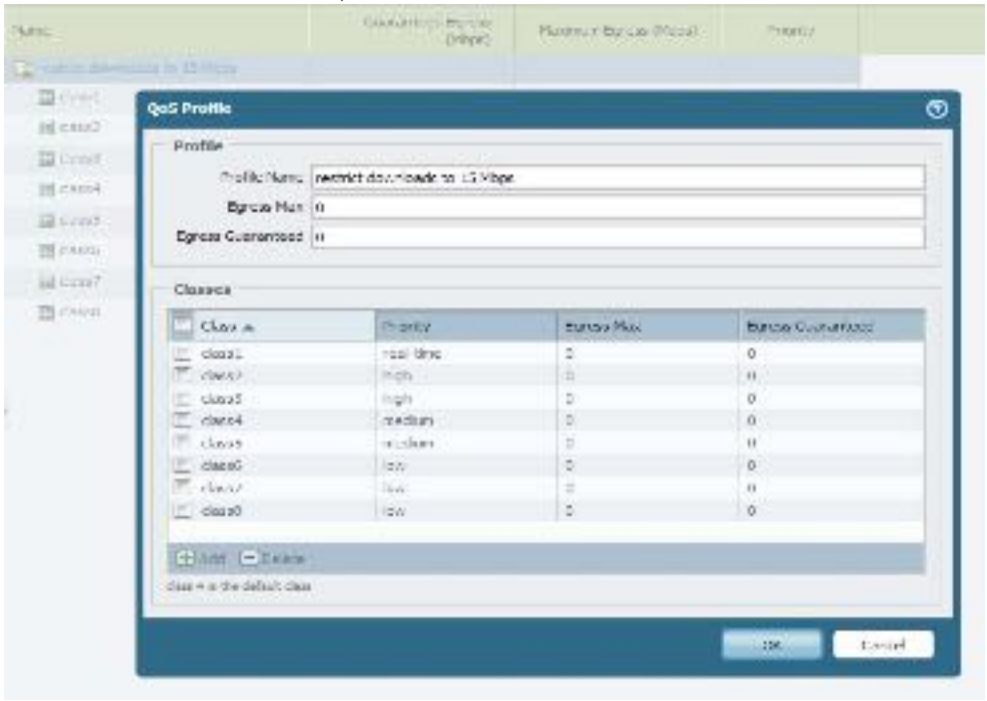


You enable QoS by applying a QoS profile to the egress interface for network, application or user traffic (or specifically, for cleartext or tunneled traffic). An interface configured with QoS shapes traffic according to the QoS profile class definitions and the traffic associated with those classes in the QoS policy.

A default QoS Profile is available on the firewall. The default profile and the classes defined in the profile do not have predefined maximum or guaranteed bandwidth limits.

You can set bandwidth limits for a QoS profile and/or set limits for individual QoS classes within the QoS profile. The total guaranteed bandwidth limits of all eight QoS classes in a QoS Profile cannot exceed the total bandwidth allocated to that QoS Profile. Enabling QoS on a physical interface includes setting the maximum bandwidth for traffic leaving the firewall through this interface. A QoS profile's guaranteed bandwidth (the **Egress Guaranteed** field) should not exceed the bandwidth allocated to the physical interface that QoS is enabled on.

For details, see Create a QoS profile.

## QoS Classes

A QoS class determines the priority and bandwidth for traffic it is assigned to. In the web interface, use the QoS profile to define QoS classes (**Network** > **Network Profiles** > **QoS Profile**):

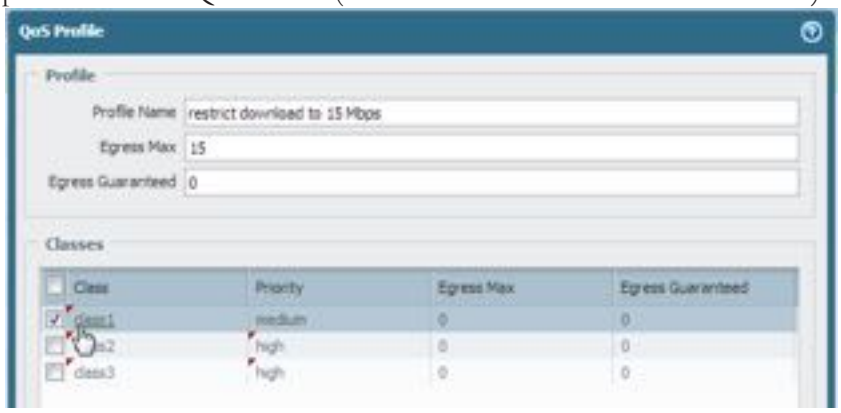


Defining a QoS class includes setting the class's Priority, maximum bandwidth (Egress Max), and guaranteed bandwidth (Egress Guaranteed).

Real-time priority is typically used for applications that are particularly sensitive to latency, such as voice and video applications.

Use the QoS policy to assign a QoS class to specified traffic (**Policies** > **QoS**):

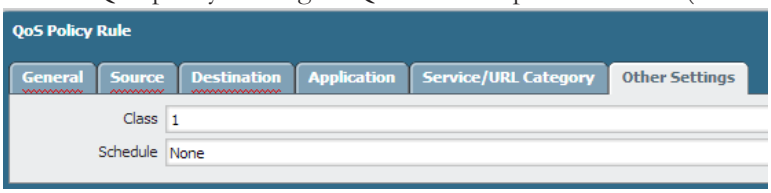


There are up to eight definable QoS classes in a single QoS profile. Unless otherwise configured, traffic that does not match a QoS class is assigned a class of 4.

 © Palo Alto Networks, Inc.

QoS priority queuing and bandwidth management, the fundamental mechanisms of a QoS configuration, are configured within the QoS class definition (see Step 3). Queuing priority is determined by the priority set for a QoS class. Bandwidth management is determined according to the maximum and guaranteed bandwidths set for a QoS class.

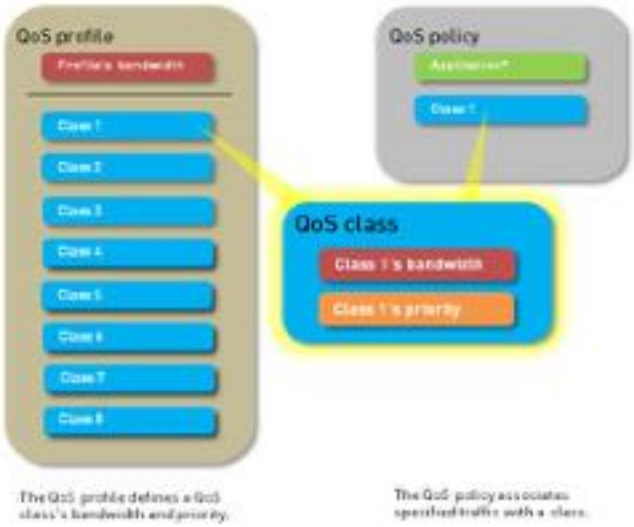


The queuing and bandwidth management mechanisms determine the order of traffic and how traffic is handled upon entering or leaving a network:

- **QoS Priority:** One of four QoS priorities can be defined in a QoS class: real-time, high, medium, and low. When a QoS class is associated with specific traffic, the priority defined in that QoS class is assigned to the traffic. Packets in the traffic flow are then queued according to their priority until the network is ready to process them. This method of priority queuing provides the capability to ensure that important traffic, applications, or users takes precedence.
- **QoS Class Bandwidth Management:** QoS class bandwidth management provides the capability to control traffic flows on a network so that traffic does not exceed network capacity, resulting in network congestion, or to allocate specific bandwidth limits to traffic, applications, or users. You can set overall limits on bandwidth using the QoS profile or set limits for individual QoS classes. A QoS profile and QoS classes in the profile have guaranteed and maximum bandwidth limits. The guaranteed bandwidth limit (Egress Guaranteed) ensures that any amount of traffic up to that set bandwidth limit is processed. The maximum bandwidth limit (Egress Max) sets the total limit of bandwidth allocated to either the QoS Profile or QoS Class. Traffic in excess of the Maximum Bandwidth limit is dropped. The total bandwidth limits and guaranteed bandwidth limits of QoS classes in a QoS profile cannot exceed the bandwidth limit of the QoS profile.

## QoS Policy

In a QoS configuration, the QoS policy identifies traffic requiring QoS treatment (either preferential treatment or bandwidth-limiting) using a defined parameter or multiple parameters and assigns it a class.

Use the QoS Policy, similar to a security policy, to set the criteria that identifies traffic:

- Applications and application groups.
- Source zones, source addresses, and source users.
- Destination zones and destination addresses.
- Services and service groups limited to specific TCP and/or UDP port numbers.
- URL categories, including custom URL categories.

The QoS Policy in the web interface (**Policies** > **QoS**) allows you to associate the criteria used to specify traffic with a QoS class.

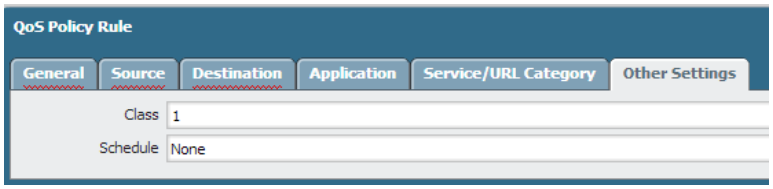


## QoS Egress Interface

Enabling a QoS profile on the egress interface of the traffic identified for QoS treatment completes a QoS configuration. The ingress interface for QoS traffic is the interface on which the traffic enters the firewall. The egress interface for QoS traffic is the interface that traffic leaves the firewall from. QoS is always enabled and enforced on a traffic flow's egress interface. The egress interface in a QoS configuration can either be the external- or internal-facing interface of the firewall, depending on the flow of the traffic receiving QoS treatment.

For example, in an enterprise network, if you are limiting employees' download traffic from a specific website, the egress interface in the QoS configuration is the firewall's internal interface, as the traffic flow is from the Internet, through the firewall, and to your company network. Alternatively, when limiting employees' upload traffic to the same website, the egress interface in the QoS configuration is the firewall's external interface, as the traffic you are limiting flows from your company network, through the firewall, and then to the Internet.

See Step 3 to learn how to Identify the egress interface for applications that you identified as needing QoS treatment.

 © Palo Alto Networks, Inc.

## QoS Cleartext and Tunneled Traffic

Within the QoS physical interface configuration, you can provide more granular QoS settings for cleartext traffic and tunneled traffic leaving through the interface. Individual tunnel interfaces can be assigned different QoS Profiles. Cleartext traffic can be assigned different QoS Profiles according to traffic's source interface and source subnet. In this case, a source interface and source subnet can be associated with a QoS Profile. If you choose not to select cleartext or tunneled traffic for unique QoS treatment, enabling QoS on an interface requires selecting a default QoS Profile to determine how to shape traffic for specific tunnel interfaces or, in the case of cleartext traffic, source interfaces and source subnets.

On Palo Alto Networks firewalls, the term "tunneled traffic" refers to tunnel interface traffic, specifically IPSec traffic in tunnel mode.

# Configure QoS

Use the following task to configure Quality of Service (QoS), including how to create a QoS profile, create a QoS policy, and enable QoS on an interface.

| Configure QoS | | |
|---|---|---|
| Step 1 | Identify traffic to apply QoS to.<br><br>This example shows how to use QoS to limit web browsing. | Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.<br><br>Click any application name to display detailed application information. |
| Step 2 | Identify the egress interface for applications that you identified as needing QoS treatment.<br><br>**Tip:** The egress interface for traffic depends on the traffic flow. If you are shaping incoming traffic, the egress interface is the internal-facing interface. If you are shaping outgoing traffic, the egress interface is the external-facing interface. | Select **Monitor > Logs > Traffic** to view the device's traffic logs.<br><br>To filter and only show logs for a specific application:<br>• If an entry is displayed for the application, click the underlined link in the Application column then click the Submit icon .<br>• If an entry is not displayed for the application, click the Add Log icon and search for the application .<br><br>The **Egress I/F** in the traffic logs displays each application's egress interface. To display the **Egress I/F** column if it is not displayed by default:<br>• Click any column header to add a column to the log:<br> 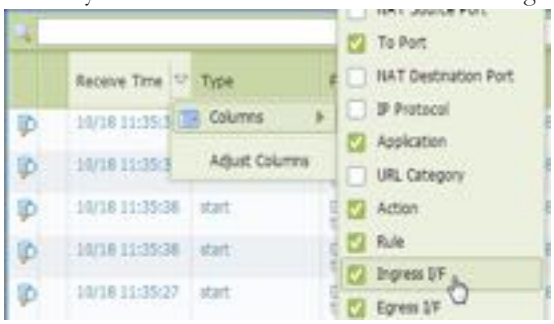 <br>• Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface listed in the Destination section:<br> <br>In this example, the egress interface for web-browsing traffic is ethernet 1/1. |

 © Palo Alto Networks, Inc.

**Configure QoS (Continued)**

Step 3 Create a QoS profile.

You can edit any existing QoS profile, including the default by clicking the QoS profile name.

1. Select **Network > Network Profiles > QoS Profile** and click **Add** to open the **QoS Profile** dialog.
2. Enter a descriptive **Profile Name**.
3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS Profile.
4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS Profile.

   Any traffic that exceeds the QoS Profile's egress guaranteed limit is best effort but is not guaranteed.
5. In the Classes section, specify how to treat up to eight individual QoS classes:
   a. Click **Add** to add a class to the QoS Profile.
   b. Select the **Priority** for the class.
   c. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.
   d. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.
6. Click **OK** to save the QoS Profile.

In the following example, the QoS Profile named Limit Web Browsing limits traffic identified as Class 2 traffic to maximum bandwidth of 50 Mbps and a guaranteed bandwidth of 2 Mbps. Any traffic that is associated with class 2 in a the QoS policy (Step 4) is subject to these limits.

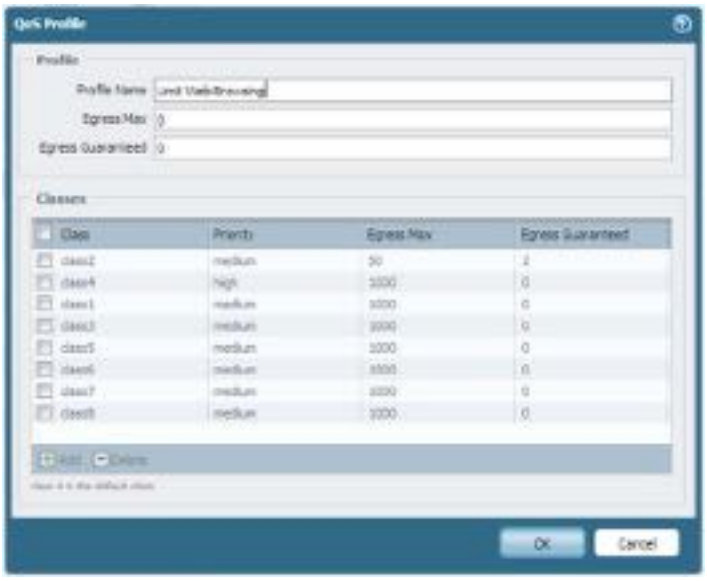

| Configure QoS (Continued) | |
|---|---|
| Step 4 Create a QoS policy. | 1. Select **Policies > QoS** and click **Add** to open the QoS Policy Rule dialog.<br>2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**.<br>3. Specify the traffic to which the QoS Policy Rule will apply. Use the **Source**, **Destination**, **Application**, and **Service/URL Category** tabs to define matching parameters for identifying traffic.<br>For example, select the **Application** tab, click **Add** and select web-browsing to apply the QoS Policy Rule to that application:<br> <br>(Optional) Define additional parameters. For example, on the **Source** tab, click **Add** to limit a specific user's web-browsing, in this case, user1:<br> 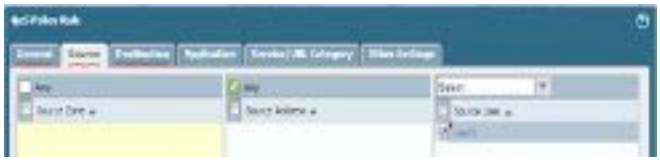 <br>4. On the **Other Settings** tab, select a QoS Class to assign to the QoS Policy Rule. For example, assign Class 2 to the user1's web-browsing traffic:<br> <br>5. Click **OK** to save the QoS Policy Rule. |


© Palo Alto Networks, Inc.

**Configure QoS (Continued)**

Step 5 Enable the QoS Profile on a physical interface.

You can configure settings to select cleartext and tunneled traffic for unique QoS treatment, in addition to the QoS configuration on the physical interface:

- To configure specific settings for cleartext traffic using traffic's source interface and source subnet as criteria for QoS identification and treatment, perform Step 5 - 4.
- To apply a QoS profile to a specific tunnel interface(s), perform Step 5 - 5

For more information, see QoS Cleartext and Tunneled Traffic.

It is a best practice to always define the **Egress Max** value for a QoS interface.

1. Select **Network > QoS** and click **Add** to open the QoS Interface dialog.
2. Enable QoS on the physical interface:
   a. On the **Physical Interface** tab, select the **Interface Name** of the interface to apply the QoS Profile to.

      In the example, Ethernet 1/1 is the egress interface for web-browsing traffic (see Step 2).

   b. Select **Turn on QoS feature on this interface**.
3. On the **Physical Interface** tab, select a QoS profile to apply by default to all **Clear Text** traffic.

   (Optional) Use the Tunnel Interface field to apply a QoS profile by default to all tunneled traffic.

For example, enable QoS on ethernet 1/1 and apply the QoS Profile named Limit Web Browsing as the default QoS Profile for clear text traffic.

4. (Optional) On the **Clear Text Traffic** tab, configure more granular QoS settings for cleartext traffic:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
   - Click **Add** to apply a QoS Profile to selected clear text traffic, further selecting the traffic for QoS treatment according to source interface and source subnet (creating a QoS node).
5. (Optional) On the **Tunneled Traffic** tab, configure more granular QoS settings for tunnel interfaces:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
   - Click **Add** to associate a selected tunnel interface with a QoS Profile.
6. Click **OK** to save the QoS Profile.
7. **Commit** the changes to enable the Qos Profile on the interface.

| Configure QoS (Continued) | |
|---|---|
| Step 6 Verify QoS configuration. | Select **Network > QoS** to view the **QoS Policies** page and click the **Statistics** link to view QoS bandwidth, active sessions of a selected QoS node or class, and active applications for the selected QoS node or class.<br>For example, see the statistics for ethernet 1/1 with QoS enabled:<br> 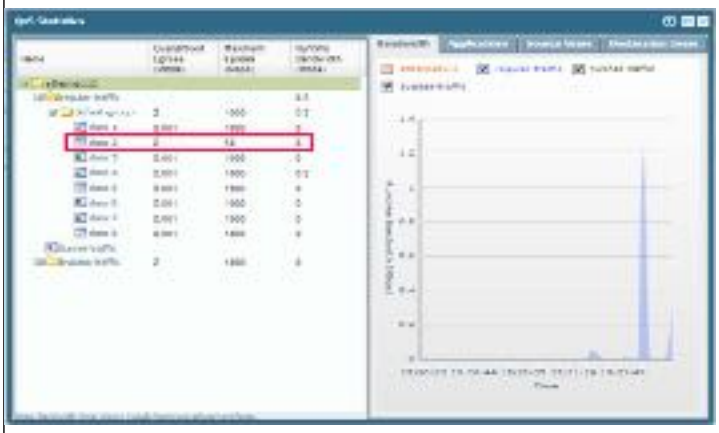 <br>Class 2 traffic limited to 2 Mbps of guaranteed bandwidth and a maximum bandwidth of 50 Mbps.<br>Continue to click the tabs to display further information regarding applications, source users, destination users, security rules and QoS rules.<br>Bandwidth limits shown on the **QoS Statistics** window include a hardware adjustment factor. |


© Palo Alto Networks, Inc.

# Configure QoS for a Virtual System

QoS can be configured for a single or several virtual systems within a Palo Alto Networks firewall. Because a virtual system is an independent firewall, QoS must be configured independently for a single virtual system to apply a QoS configuration to only that virtual system.

Configuring QoS for a virtual system is similar to configuring QoS on a physical firewall, with the exception that configuring QoS for a virtual system requires specifying the traffic flow's source and destination zones and source and destination interfaces. Because a virtual system exists without set physical boundaries (such as a physical interface) that traffic flows through, specifying source and destination zones and interfaces of a traffic flow allows you to control and shape traffic for that virtual system specifically, as a traffic flow spans more than one virtual system in a virtual environment.

The example below shows two virtual systems configured within a firewall. VSYS 1 (purple) and VSYS 2 (red) each have QoS configured to prioritize or limit two distinct traffic flows, indicated by their corresponding purple (VSYS 1) and red (VSYS 2) lines. The QoS nodes indicate the points at which QoS traffic is identified and then shaped in each virtual system.

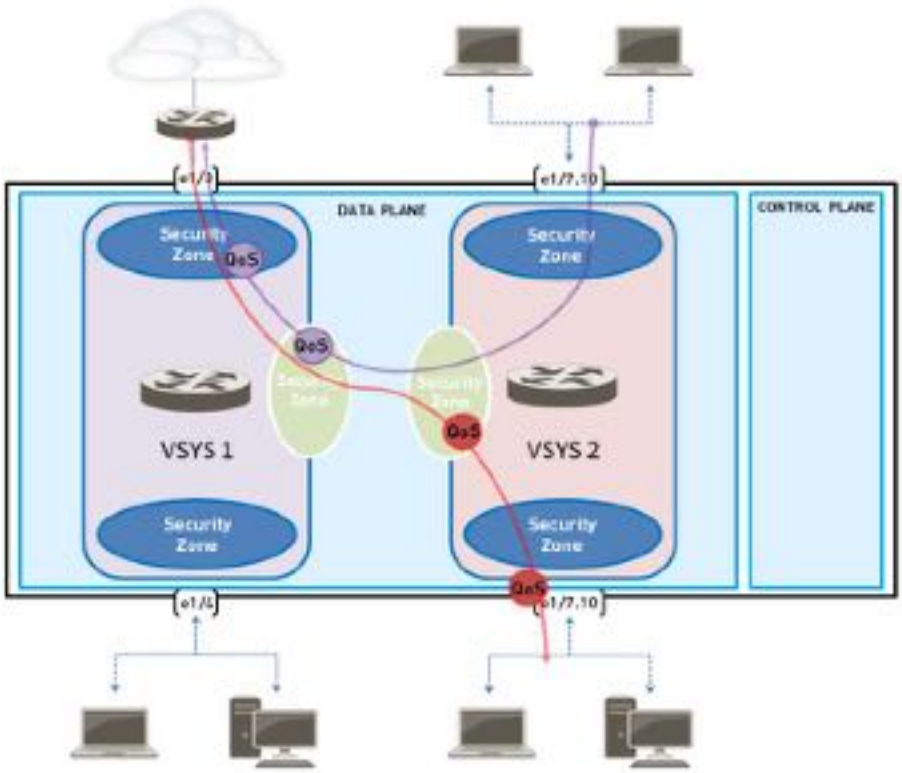


| Configure QoS in a Virtual System Environment | | |
|---|---|---|
| Step 1 | Confirm that the appropriate interfaces, virtual routers, and security zones are associated with each virtual system. | • To view configured interfaces, select **Network > Interface**.<br>• To view configured zones, select **Network > Zones**.<br>• To view information on defined virtual routers, select **Network > Virtual Routers**. |

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 2 Identify traffic to apply QoS to. | Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.<br>To view information for a specific virtual system, select the virtual system from the **Virtual System** drop-down:<br> 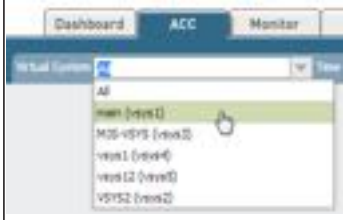 <br><br>Click any application name to display detailed application information. |

 © Palo Alto Networks, Inc.

**Configure QoS in a Virtual System Environment**

Step 3 Identify the egress interface for applications that you identified as needing QoS treatment.

In a virtual system environment, QoS is applied to traffic on the traffic's egress point on the virtual system. Depending on a virtual system's configuration and the QoS policy, the egress point of QoS traffic could be associated with a physical interface or could be a configured zone.

This example shows how to limit web-browsing traffic on vsys 1.

Select **Monitor > Logs > Traffic** to view the device's traffic logs. Each entry has the option to display columns with information necessary to configure QoS in a virtual system environment:

- virtual system
- egress interface
- ingress interface
- source zone
- destination zone

To display the a column if it is not displayed by default:

- Click any column header to add a column to the log:

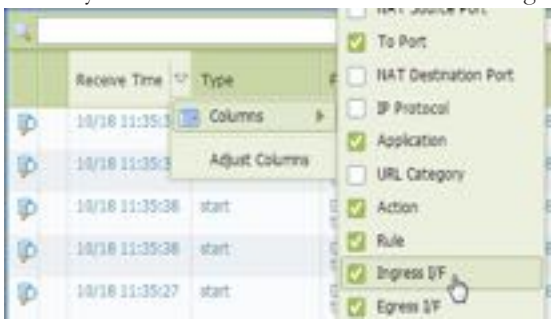


- Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface, as well as source and destination zones, in the **Source** and **Destination** sections:

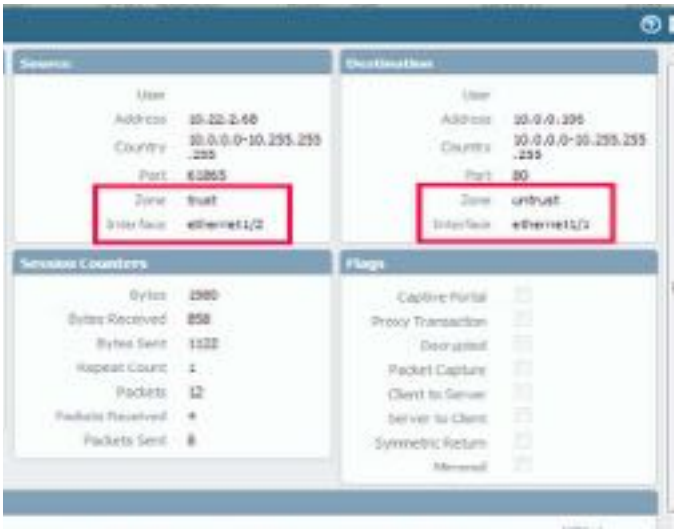

For example, for web-browsing traffic from VSYS 1, the ingress interface is ethernet 1/2, the egress interface is ethernet 1/1, the source zone is *trust* and the destination zone is *untrust*.

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 4 Create a QoS Profile.<br><br>You can edit any existing QoS Profile, including the default by clicking the profile name. | 1. Select **Network > Network Profiles > QoS Profile** and click **Add** to open the QoS Profile dialog.<br>2. Enter a descriptive **Profile Name**.<br>3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS profile.<br>4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS profile.<br><br> Any traffic that exceeds the QoS profile's egress guaranteed limit is best effort but is not guaranteed.<br>5. In the Classes section of the **QoS Profile**, specify how to treat up to eight individual QoS classes:<br>a. Click **Add** to add a class to the QoS Profile.<br>b. Select the **Priority** for the class.<br>c. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.<br>d. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.<br>6. Click **OK** to save the QoS profile. |

 © Palo Alto Networks, Inc.

**Configure QoS in a Virtual System Environment**

Step 5 Create a QoS policy.

In a multi-VSYS environment, traffic can span more than one virtual system before the ingress point on the virtual system you are configuring QoS for. Specifying source and destination zones for QoS traffic ensures that the traffic is correctly identified as it flows through the specific virtual system (in this example, vsys 1) and QoS is applied to the traffic only for that designated virtual system (and not applied to traffic for other configured virtual systems).

1. Select **Policies > QoS** and click **Add** to open the QoS Policy Rule dialog.
2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**.
3. Specify the traffic to which the QoS policy rule will apply. Use the **Source**, **Destination**, **Application**, and **Service/URL Category** tabs to define matching parameters for identifying traffic.

   For example, select the **Application** tab, click **Add** and select web-browsing to apply the QoS Policy Rule to that application:

4. On the **Source** tab, click **Add** to select vsys 1's web-browsing traffic's source zone.

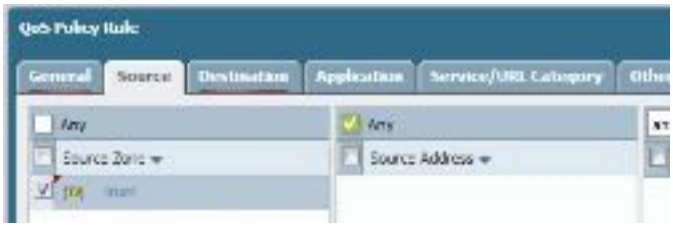


5. On the **Destination** tab, click **Add** to select the vsys 1's web-browsing traffic's destination zone.

6. On the **Other Settings** tab, and select a **QoS Class** to assign to the QoS policy rule. For example, assign Class 2 to web-browsing traffic on vsys 1:

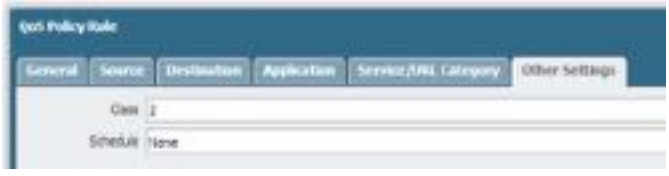


7. Click **OK** to save the QoS policy rule.

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 6 Enable the QoS Profile on a physical interface.<br><br>It is a best practice to always define the **Egress Max** value for a QoS interface. | 1. Select **Network > QoS** and click **Add** to open the QoS Interface dialog.<br>2. Enable QoS on the physical interface:<br>a. On the **Physical Interface** tab, select the **Interface Name** of the interface to apply the QoS Profile to.<br>In this example, ethernet 1/1 is the egress interface for web-browsing traffic on vsys 1 (see Step 2).<br> <br>b. Select **Turn on QoS feature on this interface**.<br>3. On the **Physical Interface** tab, select the default QoS profile to apply to all **Clear Text** traffic.<br>(Optional) Use the **Tunnel Interface** field to apply a default QoS profile to all tunneled traffic.<br>4. (Optional) On the **Clear Text Traffic** tab, configure additional QoS settings for clear text traffic:<br>• Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.<br>• Click **Add** to apply a QoS Profile to selected clear text traffic, further selecting the traffic for QoS treatment according to source interface and source subnet (creating a QoS node).<br>5. (Optional) On the **Tunneled Traffic** tab, configure additional QoS settings for tunnel interfaces:<br>• Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.<br>• Click **Add** to associate a selected tunnel interface with a QoS Profile.<br>6. Click **OK** to save changes.<br>7. **Commit** the changes. |


© Palo Alto Networks, Inc.

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 7 Verify QoS configuration. | • Select **Network > QoS** to view the QoS Policies page. The **QoS Policies** page verifies that QoS is enabled and includes a **Statistics** link. Click the Statistics link to view QoS bandwidth, active sessions of a selected QoS node or class, and active applications for the selected QoS node or class.<br>• In a multi-VSYS environment, sessions cannot span multiple systems. Multiple sessions are created for one traffic flow if the traffic passes through more than one virtual system. To browse sessions running on the firewall and view applied QoS Rules and QoS Classes, select **Monitor > Session Browser**. |

# QoS Use Case Examples

The following use cases demonstrate how to use QoS in common scenarios:

- QoS for a Single User
- QoS for Voice and Video Applications


© Palo Alto Networks, Inc.

## QoS for a Single User

A CEO finds that during periods of high network usage, she is unable to access enterprise applications respond effectively to critical business communications. The IT admin wants to ensure that all traffic to and from the CEO receives preferential treatment over other employee traffic so that she is guaranteed not only access to but high performance of critical network resources.

### Apply QoS to a Single User

Step 1 The admin creates the QoS profile *CEO_traffic* to define how traffic originating from the CEO will be treated and shaped as it flows out of the company network:

The admin assigns a guaranteed bandwidth (**Egress Guaranteed**) of 50 Mbps to ensure that the CEO will have that amount that bandwidth guaranteed to her at all times (more than she would need to use), regardless of network congestion.

The admin continues by designating Class 1 traffic as high priority and sets the profile's maximum bandwidth usage (**Egress Max**) to 1000 Mbps, the same maximum bandwidth for the interface that the admin will enable QoS on. The admin is choosing to not restrict the CEO's bandwidth usage in any way.

It is a best practice to populate the **Egress Max** field for a QoS profile, even if the max bandwidth of the profile matches the max bandwidth of the interface. The QoS profile's max bandwidth should never exceed the max bandwidth of the interface you are planning to enable QoS on.

**Apply QoS to a Single User (Continued)**

Step 2 The admin creates a QoS policy to identify the CEO's traffic (**Policies** > **QoS**) and assign it the class that he defined in the QoS profile (see Step 1). Because User-ID is configured, the admin uses the **Source** tab in the QoS policy to singularly identify the CEO's traffic by her company network username. (If User-ID is not configured, the administrator could **Add** the CEO's IP address under **Source Address**. See User-ID.):

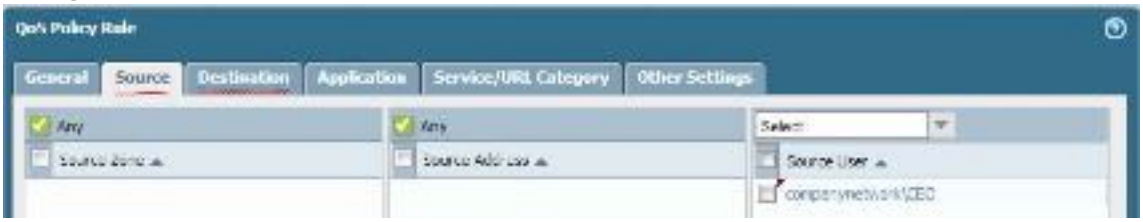


The admin associates the CEO's traffic with Class 1 (**Other Settings** tab) and then continues to populate the remaining required policy fields; the admin gives the policy a descriptive **Name** (**General** tab) and selects **Any** for the **Source Zone** (**Source** tab) and **Destination Zone** (**Destination** tab):

Step 3 Now that Class 1 is associated with the CEO's traffic, the admin enables QoS by checking **Turn on QoS feature on interface** and selecting the traffic flow's egress interface. The egress interface for the CEO's traffic flow is the external-facing interface, in this case, ethernet 1/2:

Because the admin wants to ensure that all traffic originating from the CEO is guaranteed by the QoS profile and associated QoS policy he created, he selects the *CEO_traffic* to apply to **Clear Text** traffic flowing from ethernet 1/2.

 © Palo Alto Networks, Inc.

**Apply QoS to a Single User (Continued)**

Step 4 After committing the QoS configuration, the admin navigates to the **Network > QoS** page to confirm that the QoS profile *CEO_traffic* is enabled on the external-facing interface, ethernet 1/2:

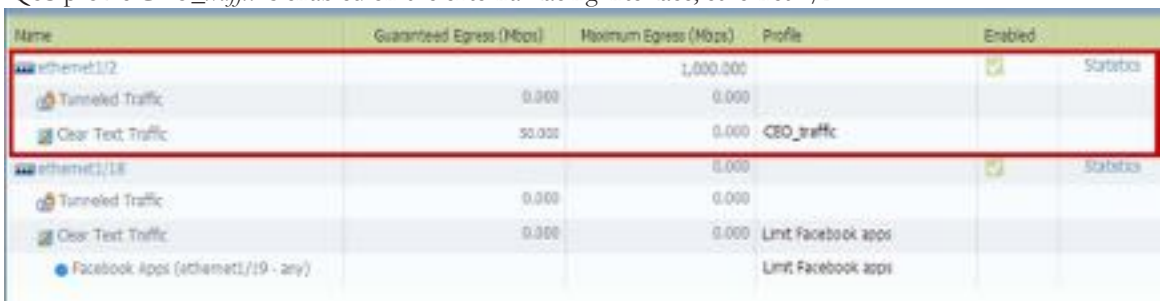


He clicks **Statistics** to view how traffic originating with the CEO (Class 1) is being shaped as it flows from ethernet 1/2:

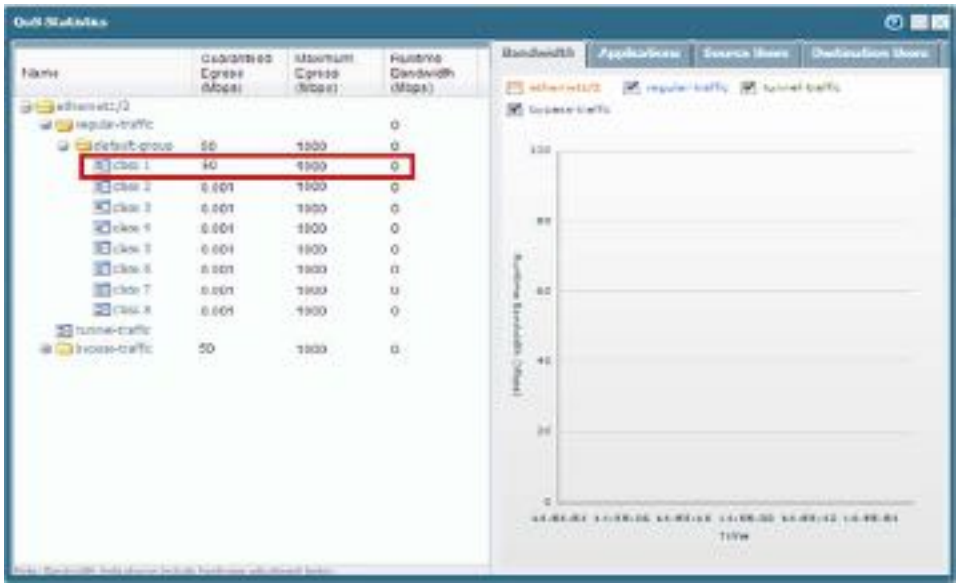

This case demonstrates how to apply QoS to traffic originating from a single source user. However, if you also wanted to guarantee or shape traffic to a destination user, you could configure a similar QoS setup. Instead of, or in addition to this work flow, create a QoS policy that specifies the user's IP address as the **Destination Address** on the **Policies > QoS** page (instead of specifying the user's source information, as shown in Step 2) and then enable QoS on the network's internal-facing interface on the **Network > QoS** page (instead of the external-facing interface, as shown in Step 3.)

## QoS for Voice and Video Applications

Voice and video traffic is particularly sensitive to measurements that the QoS feature shapes and controls, especially latency and jitter. For voice and video transmissions to be audible and clear, voice and video packets cannot be dropped, delayed, or delivered inconsistently. A best practice for voice and video applications, in addition to guaranteeing bandwidth, is to guarantee priority to voice and video traffic. In this example, employees at a company branch office are experiencing difficulties and unreliability in using video conferencing and Voice-over-IP (VoIP) technologies to conduct business communications with other branch offices, with partners, and with customers. An IT admin intends to implement QoS in order to address these issues and ensure effective and reliable business communication for the branch employees. Because the admin wants to guarantee QoS to both incoming and outgoing network traffic, he will enable QoS on both the firewall's internal- and external-facing interfaces.

**Ensure Quality for Voice and Video Applications**

Step 1 The admin creates a QoS profile, defining Class 2 so that any traffic associated with Class 2 receives real-time priority and on an interface with a maximum bandwidth of 1000 Mbps, is guaranteed a bandwidth of 250 Mbps at all times, including peak periods of network usage.

Real-time priority is typically recommended for applications affected by latency, and is particularly useful in guaranteeing performance and quality of voice and video applications.

On the **Network > Network Profiles > Qos Profile** page, the admin clicks **Add**, enters the **Profile Name** *ensure voip-video traffic* and defines Class 2 traffic.

 © Palo Alto Networks, Inc.

**Ensure Quality for Voice and Video Applications (Continued)**

Step 2 The admin creates a QoS policy to identify voice and video traffic. Because the company does not have one standard voice and video application, the admin wants to ensure QoS is applied to a few applications that are widely and regularly used by employees to communicate with other offices, with partners, and with customers. On the **Policies > QoS > QoS Policy Rule > Applications** tab, the admin clicks **Add** and opens the **Application Filter** window. The admin continues by selecting criteria to filter the applications he wants to apply QoS to, choosing the Subcategory *voip-video*, and narrowing that down by specifying only voip-video applications that are both low-risk and widely-used.

The application filter is a dynamic tool that, when used to filter applications in the QoS policy, allows QoS to be applied to all applications that meet the criteria of *voip-video*, *low risk*, and *widely used* at any given time.

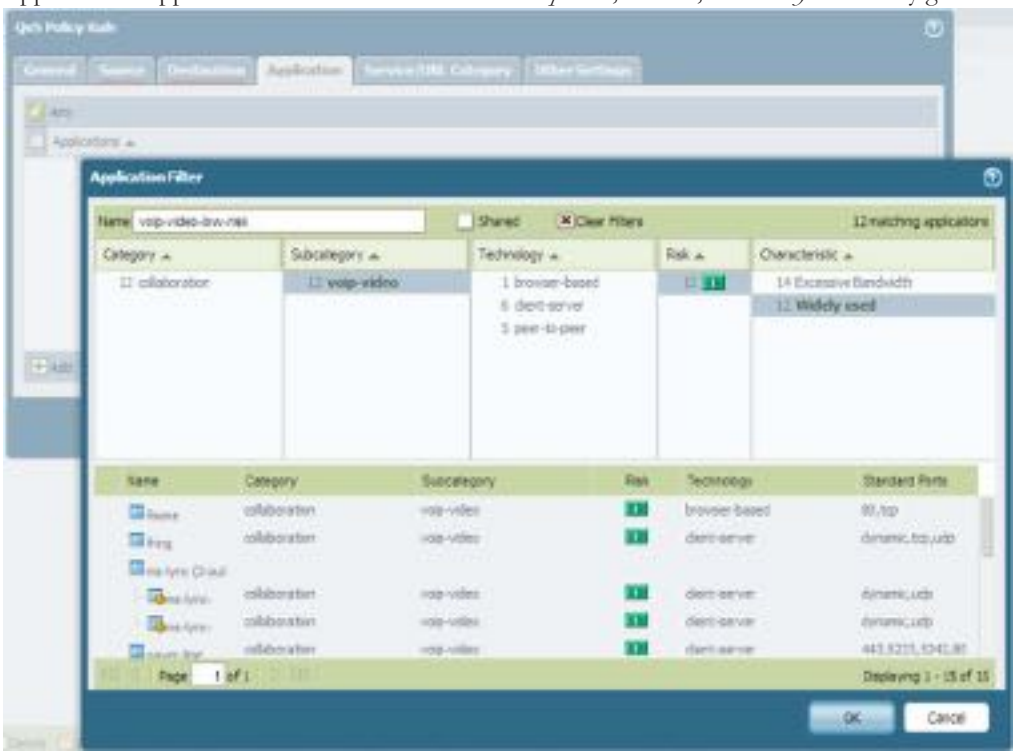


The admin names the **Application Filter** *voip-video-low-risk* and includes it in the QoS policy:

The admin names the QoS policy *Voice-Video* and associates the *voip-video-low-risk* application filter with Class 2 traffic (as he defined it in Step 1). He is going to use the *Voice-Video* QoS policy for both incoming and outgoing QoS traffic, so he sets **Source** and **Destination** information to **Any**:

**Ensure Quality for Voice and Video Applications (Continued)**

Step 3 Because the admin wants to ensure QoS for both incoming and outgoing voice and video communications, he enables QoS on the network's external-facing interface (to apply QoS to outgoing communications) and to the internal-facing interface (to apply QoS to incoming communications).

The admin begins by enabling the QoS profile he created in Step 1, *ensure voice-video traffic* (Class 1 in this profile is associated with policy created in Step 2, *Voice-Video*) on the external-facing interface, in this case, ethernet 1/2.

He then enables the same QoS profile *ensure voip-video traffic* on the internal-facing interface, in this case, ethernet 1/1.

Step 4 The admin confirms that QoS is enabled for both incoming and outgoing voice and video traffic:

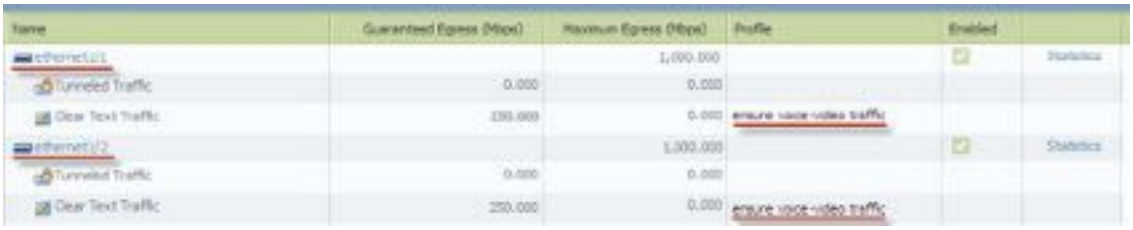


The admin has successfully enabled QoS on both the network's internal- and external-facing interfaces. Real-time priority is now ensured for voice and video application traffic as it flows both into and out of the network, ensuring that these communications, which are particularly sensitive to latency and jitter, can be used reliably and effectively to perform both internal and external business communications.

 © Palo Alto Networks, Inc.

# VPNs

Virtual private networks (VPNs) create tunnels that allow users/systems to connect securely over a public network, as if they were connecting over a local area network (LAN). To set up a VPN tunnel, you need a pair of devices that can authenticate each other and encrypt the flow of information between them. The devices can be a pair of Palo Alto Networks firewalls, or a Palo Alto Networks firewall along with a VPN-capable device from another vendor.


- VPN Deployments
- Site-to-Site VPN Overview
- Site-to-Site VPN Concepts
- Set Up Site-to-Site VPN
- Site-to-Site VPN Quick Configs


© Palo Alto Networks, Inc.

# VPN Deployments

The Palo Alto Networks firewall supports the following VPN deployments:

- **Site-to-Site VPN**— A simple VPN that connects a central site and a remote site, or a hub and spoke VPN that connects a central site with multiple remote sites. The firewall uses the IP Security (IPSec) set of protocols to set up a secure tunnel for the traffic between the two sites. See Site-to-Site VPN Overview.
- **Remote User-to-Site VPN**—A solution that uses the GlobalProtect agent to allow a remote user to establish a secure connection through the firewall. This solution uses SSL and IPSec to establish a secure connection between the user and the site. Refer to the GlobalProtect Administrator's Guide.
- **Large Scale VPN**— The Palo Alto Networks GlobalProtect Large Scale VPN (LSVPN) provides a simplified mechanism to roll out a scalable hub and spoke VPN with up to 1024 satellite offices. The solution requires Palo Alto Networks firewalls to be deployed at the hub and at every spoke. It uses certificates for device authentication, SSL for securing communication between all components, and IPSec to secure data. See Large Scale VPN (LSVPN).

**Figure: VPN Deployments**

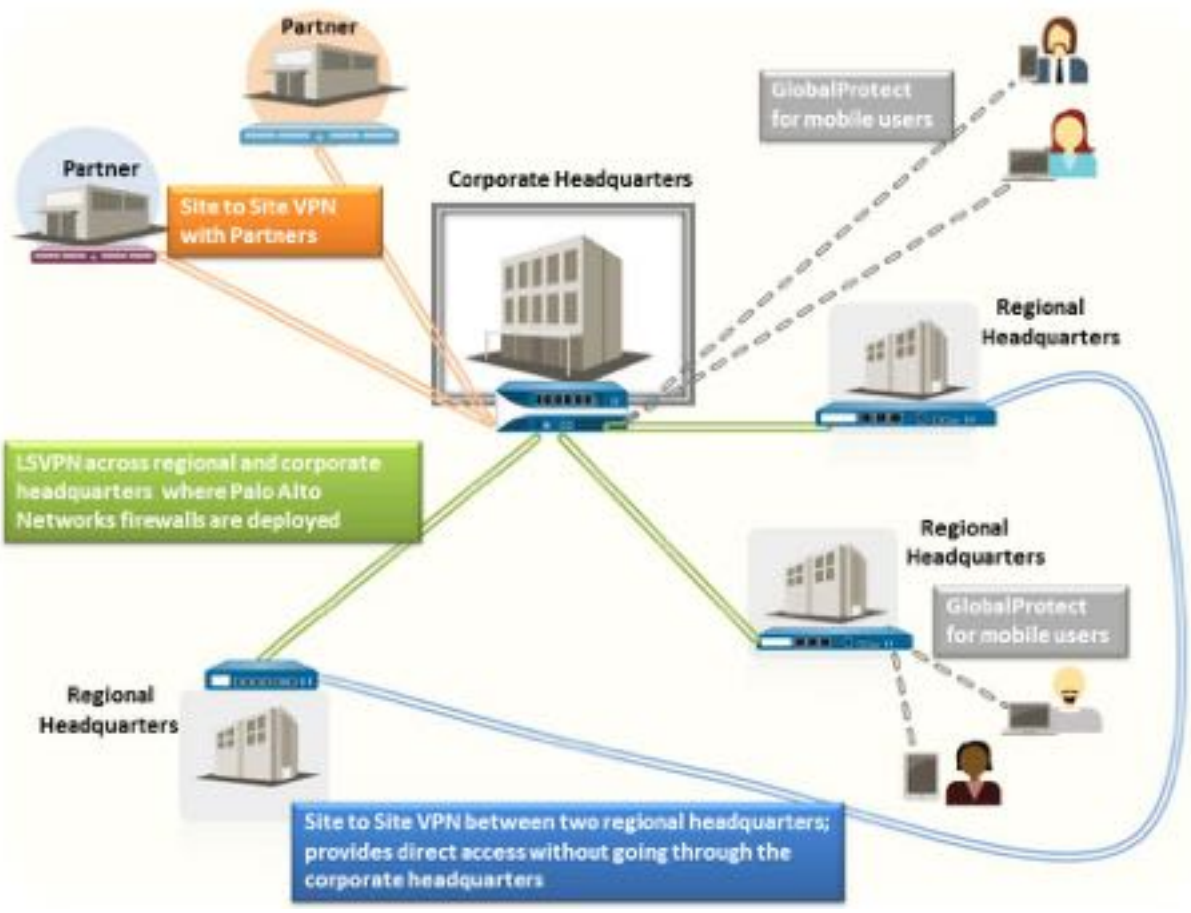


 © Palo Alto Networks, Inc.

# Site-to-Site VPN Overview

A VPN connection that allows you to connect two Local Area Networks (LANs) is called a site-to-site VPN. You can configure route-based VPNs to connect Palo Alto Networks firewalls located at two sites or to connect a Palo Alto Networks firewall with a third-party security device at another location. The firewall can also interoperate with third-party policy-based VPN devices; the Palo Alto Networks firewall supports route-based VPN.

The Palo Alto Networks firewall sets up a route-based VPN, where the firewall makes a routing decision based on the destination IP address. If traffic is routed to a specific destination through a VPN tunnel, then it is handled as VPN traffic.

The IP Security (IPSec) set of protocols is used to set up a secure tunnel for the VPN traffic, and the information in the TCP/IP packet is secured (and encrypted if the tunnel type is ESP). The IP packet (header and payload) is embedded in another IP payload, and a new header is applied and then sent through the IPSec tunnel. The source IP address in the new header is that of the local VPN peer and the destination IP address is that of the VPN peer on the far end of the tunnel. When the packet reaches the remote VPN peer (the firewall at the far end of the tunnel), the outer header is removed and the original packet is sent to its destination.

In order to set up the VPN tunnel, first the peers need to be authenticated. After successful authentication, the peers negotiate the encryption mechanism and algorithms to secure the communication. The Internet Key Exchange (IKE) process is used to authenticate the VPN peers, and IPSec Security Associations (SAs) are defined at each end of the tunnel to secure the VPN communication. IKE uses digital certificates or preshared keys, and the Diffie Hellman keys to set up the SAs for the IPSec tunnel.The SAs specify all of the parameters that are required for secure transmission— including the security parameter index (SPI), security protocol, cryptographic keys, and the destination IP address—encryption, data authentication, data integrity, and endpoint authentication.

The following figure shows a VPN tunnel between two sites. When a client that is secured by VPN Peer A needs content from a server located at the other site, VPN Peer A initiates a connection request to VPN Peer B. If the security policy permits the connection, VPN Peer A uses the IKE Crypto profile parameters (IKE phase 1) to establish a secure connection and authenticate VPN Peer B. Then, VPN Peer A establishes the VPN tunnel using the IPSec Crypto profile, which defines the IKE phase 2 parameters to allow the secure transfer of data between the two sites.

**Figure: Site-to-Site VPN**

# Site-to-Site VPN Concepts

A VPN connection provides secure access to information between two or more sites. In order to provide secure access to resources and reliable connectivity, a VPN connection needs the following components:

- IKE Gateway
- Tunnel Interface
- Tunnel Monitoring
- Internet Key Exchange (IKE) for VPN

## IKE Gateway

The Palo Alto Networks firewalls or a firewall and another security device that initiate and terminate VPN connections across the two networks are called the IKE Gateways. To set up the VPN tunnel and send traffic between the IKE Gateways, each peer must have an IP address—static or dynamic—or FQDN. The VPN peers use preshared keys or certificates to mutually authenticate each other.

The peers must also negotiate the mode—main or aggressive—for setting up the VPN tunnel and the SA lifetime in IKE Phase 1. Main mode protects the identity of the peers and is more secure because more packets are exchanged when setting up the tunnel. Main mode is the recommended mode for IKE negotiation if both peers support it. Aggressive mode uses fewer packets to set up the VPN tunnel and is hence faster but a less secure option for setting up the VPN tunnel.

See Set up an IKE Gateway for configuration details.

## Tunnel Interface

To set up a VPN tunnel, the Layer 3 interface at each end must have a logical *tunnel* interface for the firewall to connect to and establish a VPN tunnel. A tunnel interface is a logical (virtual) interface that is used to deliver traffic between two endpoints. Each tunnel interface can have a maximum of 10 IPSec tunnels; this means that up to 10 networks can be associated with the same tunnel interface on the firewall.

The tunnel interface must belong to a security zone to apply policy and it must be assigned to a virtual router in order to use the existing routing infrastructure. Ensure that the tunnel interface and the physical interface are assigned to the same virtual router so that the firewall can perform a route lookup and determine the appropriate tunnel to use.

Typically, the Layer 3 interface that the tunnel interface is attached to belongs to an external zone, for example the untrust zone. While the tunnel interface can either be in the same security zone as the physical interface, for added security and better visibility, you can create a separate zone for the tunnel interface. If you create a separate zone for the tunnel interface, say a VPN zone, you will need to create security policies to enable traffic to flow between the VPN zone and the trust zone.

 © Palo Alto Networks, Inc.

To route traffic between the sites, a tunnel interface does not require an IP address. An IP address is only required if you want to enable tunnel monitoring or if you are using a dynamic routing protocol to route traffic across the tunnel. With dynamic routing, the tunnel IP address serves as the next hop IP address for routing traffic to the VPN tunnel.

If you are configuring the Palo Alto Networks firewall with a VPN peer that performs policy-based VPN, you must configure a local and remote Proxy ID when setting up the IPSec tunnel. Each peer compares the Proxy-IDs configured on it with what is actually received in the packet in order to allow a successful IKE phase 2 negotiation. If multiple tunnels are required, configure unique Proxy IDs for each tunnel interface; a tunnel interface can have a maximum of 250 ProxyIDs. Each Proxy ID counts towards the IPSec VPN tunnel capacity of the firewall, and the tunnel capacity varies by the firewall model.

See Set up an IPSec Tunnel for configuration details.

## Tunnel Monitoring

For a VPN tunnel, you can check connectivity to a destination IP address across the tunnel. The network monitoring profile on the firewall allows you to verify connectivity (using ICMP) to a destination IP address or a next hop at a specified polling interval, and to specify an action on failure to access the monitored IP address.

If the destination IP is unreachable, you either configure the firewall to wait for the tunnel to recover or configure automatic failover to another tunnel. In either case, the firewall generates a system log that alerts you to a tunnel failure and renegotiates the IPSec keys to accelerate recovery.

The default monitoring profile is configured to wait for the tunnel to recover; the polling interval is 3 seconds and the failure threshold is 5.

See Set up Tunnel Monitoring for configuration details.

## Internet Key Exchange (IKE) for VPN

The IKE process allows the VPN peers at both ends of the tunnel to encrypt and decrypt packets using mutually agreed upon keys or certificate and method of encryption. The IKE process occurs in two phases: IKE Phase 1 and IKE Phase 2. Each of these phases use keys and encryption algorithms that are defined using cryptographic profiles— IKE crypto profile and IPSec crypto profile—and the result of the IKE negotiation is a Security Association (SA). An SA is a set of mutually agreed upon keys and algorithms that will be used by both VPN peers to allow the flow of data across the VPN tunnel. The following illustration depicts the key exchange process for setting up the VPN tunnel:

© Palo Alto Networks, Inc.

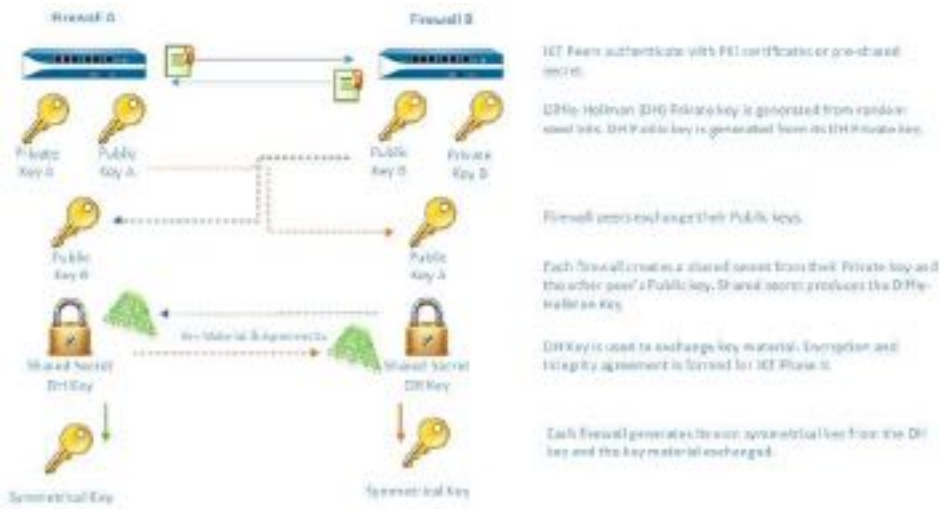

## IKE Phase 1

In this phase, the firewalls use the parameters defined in the IKE Gateway configuration and the IKE Crypto profile to authenticate each other and set up a secure control channel. IKE Phase supports the use of preshared keys or digital certificates (which use public key infrastructure, PKI) for mutual authentication of the VPN peers. Preshared keys are a simple solution for securing smaller networks because they do not require the support of a PKI infrastructure. Digital certificates can be more convenient for larger networks or implementations that require stronger authentication security.

When using certificates, make sure that the CA issuing the certificate is trusted by both gateway peers and that the maximum length of certificates in the certificate chain is 5 or less. With IKE fragmentation enabled, the firewall can reassemble IKE messages with up to 5 certificates in the certificate chain and successfully establish a VPN tunnel.

The IKE-crypto profile defines the following options that are used in the IKE SA negotiation:

- Diffie-Hellman (DH) Group for generating symmetrical keys for IKE. The Diffie Hellman algorithm uses the private key of one party and the public key of the other to create a shared secret, which is an encrypted key that is shared by both VPN tunnel peers. The DH groups supported on the firewall are: Group 1—768 bits; Group 2—1024 bits (the default); Group 5—1536 bits; Group 14—2048 bits.
- Authentication options—sha1; sha 256; sha 384; sha 512; md5
- Encryption algorithms—3des; aes128; aes192; aes256

## IKE Phase 2

After the tunnel is secured and authenticated, in phase 2 the channel is further secured for the transfer of data between the networks. IKE phase 2 uses the keys that were established in Phase 1 of the process and the IPSec Crypto profile, which defines the IPSec protocols and keys used for the SA in IKE Phase 2.

The IPSEC uses the following protocols to enable secure communication:

 © Palo Alto Networks, Inc.

- Encapsulating Security Payload (ESP)—Allows you to encrypt the entire IP packet, and authenticate the source and verify integrity of the data. While ESP requires that you encrypt and authenticate the packet, you can choose to only encrypt or only authenticate by setting the encryption option to Null; using encryption without authentication is discouraged.
- Authentication Header (AH)—Authenticates the source of the packet and verifies data integrity. AH does not encrypt the data payload and is unsuited for deployments where data privacy is important. AH is commonly used when the main concern is to verify the legitimacy of the peer, and data privacy is not required.

**Table: Algorithms Supported for IPSEC Authentication and Encryption**

| ESP | AH |
|---|---|
| Diffie Hellman Exchange options supported | |
| • Group 1—768 bits<br>• Group 2—1024 bits (the default)<br>• Group 5—1536 bits<br>• Group 14—2048 bits.<br>• no-pfs—By default, Perfect Forward Secrecy (pfs) is enabled. With PFS is enabled, a new DH key is generated in IKE phase 2 using one of the groups listed above; this key is independent of the keys exchanged in IKE phase1, and therefore allows for more secure transfer of data.<br>No-pfs implies that the DH key created at phase 1 is not renewed and a single key is used for the IPSEC SA negotiations. Both VPN peers must be enabled or disabled for Perfect Forward Secrecy. | |
| Encryption algorithms supported | |
| • 3des | |
| • aes128 | |
| • aes192 | |
| • aes256 | |
| • aes128ccm16 | |
| • null | |
| Authentication algorithms supported | |
| • md5 | • md5 |
| • sha 1 | • sha 1 |
| • sha 256 | • sha 256 |
| • sha 384 | • sha 384 |
| • sha512 | • sha 512 |
| • none | |

## Methods of Securing IPSec VPN Tunnels (IKE Phase 2)

IPSec VPN tunnels can be secured using manual keys or auto keys. In addition, IPSec configuration options include Diffie-Hellman Group for key agreement, and/or an encryption algorithm and a hash for message authentication.

- **Manual Key**—Manual key is typically used if the Palo Alto Networks firewall is establishing a VPN tunnel with a legacy device, or if you want to reduce the overhead of generating session keys. If using manual keys, the same key must be configured on both peers.

  Manual keys are not recommended for establishing a VPN tunnel because the session keys can be compromised when relaying the key information between the peers; if the keys are compromised, the data transfer is no longer secure.

- **Auto Key**— Auto Key allows you to automatically generate keys for setting up and maintaining the IPSec tunnel based on the algorithms defined in the IPSec Crypto profile.

 © Palo Alto Networks, Inc.

# Set Up Site-to-Site VPN

To set up site-to-site VPN:

- Make sure that your Ethernet interfaces, virtual routers, and zones are configured properly. For more information, see Set Up Interfaces and Zones.
- Create your tunnel interfaces. Ideally, put the tunnel interfaces in a separate zone, so that tunneled traffic can use different policies.
- Set up static routes or assign routing protocols to redirect traffic to the VPN tunnels. To support dynamic routing (OSPF, BGP, RIP are supported), you must assign an IP address to the tunnel interface.
- Define IKE gateways for establishing communication between the peers across each end of the VPN tunnel; also define the cryptographic profile that specifies the protocols and algorithms for identification, authentication, and encryption to be used for setting up VPN tunnels in IKEv1 Phase 1. See Set up an IKE Gateway and Define IKE Crypto Profiles.
- Configure the parameters that are needed to establish the IPSec connection for transfer of data across the VPN tunnel; See Set up an IPSec Tunnel. For IKEv1 Phase-2, see Define IPSec Crypto Profiles.
- (Optional) Specify how the firewall will monitor the IPSec tunnels. See Set up Tunnel Monitoring.
- Define security policies to filter and inspect the traffic.

If there is a deny rule at the end of the security rulebase, intra-zone traffic is blocked unless otherwise allowed. Rules to allow IKE and IPSec applications must be explicitly included above the deny rule.

When these tasks are complete, the tunnel is ready for use. Traffic destined for the zones/addresses defined in policy is automatically routed properly based on the destination route in the routing table, and handled as VPN traffic. For a few examples on site-to-site VPN, see Site-to-Site VPN Quick Configs.

## Set up an IKE Gateway

To set up a VPN tunnel, the VPN peers or gateways must authenticate each other using preshared keys or digital certificates and establish a secure channel in which to negotiate the IPSec security association (SA) that will be used to secure traffic between the hosts on each side.

| Set up an IKE Gateway | | |
|---|---|---|
| Step 1 | Define the new IKE Gateway. | 1. Select **Network > Network Profiles > IKE Gateway** and enter a **Name f**or the new gateway configuration.<br>2. Select the outgoing **Interface** on the firewall.<br>3. From the **Local IP address** drop down list, select the IP address that will be used to as the endpoint for the VPN connection. This is the external facing interface with a publicly routable IP address on the firewall. |
| Step 2 | Define the settings for the peer at the far end of the tunnel. | 1. Select whether the peer uses a **Static** or **Dynamic** IP address in **Peer IP Type**.<br>2. If the **Peer IP Address** is static, enter the IP address of the peer. |
| Step 3 | Select the peer authentication method.<br>This is required for static and dynamic peers. | • For configuring a **Pre-Shared Key**, see Step 4.<br>• For configuring digital **Certificates**, see Step 5. |
| Step 4 | Configure a pre-shared key. | 1. Enter a security key to use for authentication across the tunnel. This key must be the same on both peers.<br>Generate a key that is hard to crack with dictionary attacks; use a pre-shared key generator, if necessary.<br>2. Continue with Step 6. |
| Step 5 | Configure certificate-based authentication.<br>The pre-requisites for certificate-based authentication are as follows:<br>– **Obtain a signed certificate:** See Generate a Certificate on the Firewall or Obtain a Certificate from an External CA.<br>– **Configure the certificate profile**: The certificate profile provides the settings that the IKE gateway uses for negotiating and validating certificate authentication with its peer. See Configure a Certificate Profile. | 1. Select **Certificate** for the **Authentication** method and select the signed certificate from the **Local Certificate** drop-down.<br>If your device is enabled for multi virtual systems, if the certificate belongs to a virtual system, it must be in the same virtual system as the interface used for the IKE gateway.<br>2. From the **Local Identification** drop down list, choose one of the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), Distinguished Name (subject).<br>3. From the **Peer Identification** drop-down, choose one of the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), Distinguished Name (subject).<br>4. Select a **Certificate Profile** to use.<br>5. Continue with Step 6. |

 © Palo Alto Networks, Inc.

| Set up an IKE Gateway | | |
|---|---|---|
| Step 6 | Configure the additional parameters for IKE phase 1 negotiations—Exchange mode, Crypto profile, IKE fragmentation, Dead Peer Detection. | 1. Select **Network > Network Profiles > IKE Gateways** and select the **Advanced Phase 1 Options** tab.<br>2. Choose **auto**, **aggressive**, or **main** for the **Exchange Mode**.<br>When a device is set to use the **auto** exchange mode, it can accept both main mode and aggressive mode negotiation requests; however, whenever possible, it initiates negotiation and allows exchanges in main mode.<br> If the exchange mode is not set to auto, you must configure both VPN peers with the same exchange mode to allow each peer to accept negotiation requests.<br>3. Select an existing profile or keep the default profile from **IKE Crypto Profile** drop-down. For details on defining an IKE Crypto profile, see Define IKE Crypto Profiles.<br>4. Select **Passive Mode** if you want the firewall to only respond to IKE connections and never initiate them.<br>5. Select **NAT Traversal Select** to have UDP encapsulation used on IKE and UDP protocols, enabling them to pass through intermediate NAT devices.<br>6. (Only if using certificate-based authentication and the exchange mode is not set as aggressive mode) Select **Enable Fragmentation** to enable the firewall to operate with IKE Fragmentation.<br>7. Select the **Dead Peer Detection** check box and enter an **Interval** (2 - 100 seconds); For **Retry,** define the time to delay (2 - 100 seconds) before attempting to re-check availability.<br>Dead peer detection identifies inactive or unavailable IKE peers by sending an IKE phase 1 notification payload to the peer and waiting for an acknowledgment. |
| Step 7 | Save the changes. | Click **OK** and **Commit**. |

# Define Cryptographic Profiles

A cryptographic profile specifies the ciphers used for authentication and/or encryption between two IKE peers, and the lifetime of the key. The time period between each renegotiation is known as the lifetime; when the specified time expires, the firewall renegotiates a new set of keys.

For securing communication across the VPN tunnel, the firewall requires IKE and IPSec cryptographic profiles for completing IKE phase 1 and phase 2 negotiations, respectively. The firewall includes a *default* IKE crypto profile and a *default* IPSec crypto profile that is ready for use.

- Define IKE Crypto Profiles
- Define IPSec Crypto Profiles

## Define IKE Crypto Profiles

The IKE crypto profile is used to set up the encryption and authentication algorithms used for the key exchange process in IKE Phase 1, and lifetime of the keys which specifies how long the keys are valid. To invoke the profile, you must attach it to the IKE Gateway configuration.

All IKE gateways configured on the same interface or local IP address must use the same crypto profile.

| Define an IKE Crypto Profile | | |
|---|---|---|
| Step 1 | Create a new IKE profile. | 1. Select **Network > Network Profiles > IKE Crypto** and select **Add**.<br>2. Enter a **Name f**or the new profile. |
| Step 2 | Select the DH Group to use for setting up the key exchange. | Click **Add** and select the key strength that you want to use for the DH Group.<br>Select more than one DH group if you are not certain of what is supported by the VPN peer. Prioritize the list of ciphers by strength so that the strongest cipher is used for setting up the tunnel. |
| Step 3 | Select the authentication and encryption algorithm. | Click **Add** and select the **Authentication** and **Encryption** algorithms that you want to use for communication between the IKE peers.<br>Selecting multiple algorithms allows the peers to use the strongest cipher/algorithm that is supported on both IKE peers. |
| Step 4 | Specify the duration for which the key is valid. | Select the **Lifetime** for which the key is valid. The time period between each renegotiation is known as the lifetime; when the specified time expires, the firewall will renegotiate a new set of keys. |
| Step 5 | Save your IKE Crypto profile. | Click **OK** and click **Commit**. |
| Step 6 | Attach the IKE Crypto profile to the IKE Gateway configuration. | See Step 6 in Set up an IKE Gateway. |

 © Palo Alto Networks, Inc.

## Define IPSec Crypto Profiles

The IPSec crypto profile is invoked in IKE Phase 2. It specifies how the data is secured within the tunnel when Auto Key IKE is used to automatically generate keys for the IKE SAs.

| Define the IPSec Crypto Profile | | |
|---|---|---|
| Step 1 | Create a new IPSec profile. | 1. Select **Network > Network Profiles > IPSec Crypto** and select **Add**.<br>2. Enter a **Name f**or the new profile.<br>3. Select the **IPSec Protocol**—ESP or AH—that you want to apply to secure the data as it traverses across the tunnel.<br>4. Click **Add** and select the **Authentication** and **Encryption** algorithms for ESP, and **Authentication** algorithms for AH, so that the IKE peers can negotiate the keys for the secure transfer of data across the tunnel.<br>Selecting multiple algorithms allows the peers to use the strongest cipher/algorithm that is supported on both IKE peers. |
| Step 2 | Select the DH Group to use for the IPSec SA negotiations in IKE phase 2. | 1. Select the key strength that you want to use from the **DH Group** drop down.<br>Select more than one DH group if you are not certain of what key strength is supported by the peer at the other end. The strongest cipher will be used for setting up the tunnel.<br>2. Select **no-pfs**, if you do not want to renew the key that was created at phase 1, the current key is reused for the IPSEC SA negotiations. |
| Step 3 | Specify the duration of the key— time and volume of traffic. | Using a combination of time and traffic volume allows you to ensure safety of data.<br>Select the **Lifetime** or time period for which the key is valid. When the specified time expires, the firewall will renegotiate a new set of keys.<br>Select the **Lifesize** or volume of data after which the keys must be renegotiated. |
| Step 4 | Save your IPSec profile. | Click **OK** and click **Commit**. |
| Step 5 | Attach the IPSec Profile to an IPSec tunnel configuration. | See Step 4 in |

## Set up an IPSec Tunnel

The IPSec tunnel configuration allows you to authenticate and/or encrypt the data (IP packet) as it traverses across the tunnel.

If you are setting up the Palo Alto Networks firewall to work with a peer that supports policy-based VPN, you must define Proxy IDs. Devices that support policy-based VPN, use specific security rules/policies or access-lists (source addresses, destination addresses and ports) for permitting interesting traffic through an IPSec tunnel. These rules are referenced during quick mode/IKE phase 2 negotiation, and are exchanged as Proxy-IDs in the first or the second message of the process. So, if you are configuring the Palo Alto Networks firewall to work with a policy-based VPN peer, for a successful phase 2 negotiation you must define the Proxy-ID so that the setting on both peers is identical. If the Proxy-ID is not configured, because the Palo Alto Networks firewall supports route-based VPN, the default values used as Proxy-ID are source ip: 0.0.0.0/0, destination ip: 0.0.0.0/0 and application: any; and when these values are exchanged with the peer, it results in a failure to set up the VPN connection.

**Set up an IPSec Tunnel**

Step 1 Select **Network > IPSec Tunnels> General** and enter a **Name** for the new tunnel.

Step 2 Select the **Tunnel interface** that will be used to set up the IPSec tunnel.

– To create a new tunnel interface:

1. Select **Network > Interfaces > Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:

   • To use your trust zone as the termination point for the tunnel, select the zone from the drop-down. Associating the tunnel interface with the same zone (and virtual router) as the external-facing interface on which the packets enter the firewall, mitigates the need to create inter-zone routing.

   • (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example vpn-corp), and click **OK**.
4. In the **Virtual Router** drop-down, select **default**.
5. (Optional) If you want to assign an IPv4 address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.31.32.1/32.
6. If you want to assign an IPv6 address to the tunnel interface, see Step 3.
7. To save the interface configuration, click **OK**.

 © Palo Alto Networks, Inc.

| Set up an IPSec Tunnel | |
|---|---|
| Step 3 (Optional) Enable IPv6 on the tunnel interface. | 1. Select the IPv6 tab on **Network > Interfaces > Tunnel > IPv6**.<br>2. Select the check box to **Enable IPv6 on the interface**.<br>This option allows you to route IPv6 traffic over an IPv4 IPSec tunnel and will provide confidentiality between IPv6 networks. The IPv6 traffic is encapsulated by IPv4 and then ESP. To route IPv6 traffic to the tunnel, you can use a static route to the tunnel, or use OSPFv3, or use a Policy Based Forwarding (PBF) rule to direct traffic to the tunnel.<br>3. Enter the 64-bit extended unique **Interface ID** in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. By default, the firewall will use the EUI-64 generated from the physical interface's MAC address.<br>4. To enter an IPv6 **Address**, click **Add** and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. If Prefix is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.<br>a. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address.<br>b. Select **Anycast** to include routing through the nearest node. |
| Step 4 Select the type of key that will be used to secure the IPSec tunnel. | Continue to one of the following steps, depending on why type of key exchange you are using:<br>• Set up Auto Key exchange.<br>• Set up Auto Key exchange. |
| • Set up Auto Key exchange. | 1. Select the IKE Gateway. To set up an IKE gateway, see Set up an IKE Gateway.<br>2. (Optional) Select the default IPSec Crypto Profile. To create a new IPSec Profile, see Define IPSec Crypto Profiles. |

| Set up an IPSec Tunnel | |
|---|---|
| • Set up a Manual Key exchange. | 1. Set up the parameters for the local firewall:<br>a. Specify the **SPI** for the local firewall. SPI is a 32-bit hexadecimal index that is added to the header for IPSec tunneling to assist in differentiating between IPSec traffic flows; it is used to create the SA required for establishing a VPN tunnel.<br>b. Select the **Interface** that will be the tunnel endpoint, and optionally select the IP address for the local interface that is the endpoint of the tunnel.<br>c. Select the protocol to be used—**AH** or **ESP**.<br>d. For AH, select the **Authentication** method from the drop-down and enter a **Key** and then **Confirm Key**.<br>e. For ESP, select the **Authentication** method from the drop-down and enter a **Key** and then **Confirm Key**. Then, select the **Encryption** method and enter a **Key** and then **Confirm Key**, if needed.<br>2. Set up the parameters that pertain to the remote VPN peer.<br>a. Specify the **SPI** for the remote peer.<br>b. Enter the **Remote Address**, the IP address of the remote peer. |
| Step 5 Protect against a replay attack.<br>A replay attack occurs when a packet is maliciously intercepted and retransmitted by the interceptor. | Select the **Show Advanced Options** check box, select **Enable Replay Protection** to detect and neutralize against replay attacks. |
| Step 6 Preserve the Type of Service header for the priority or treatment of IP packets. | In the **Show Advanced Options** section, select **Copy TOS Header**. This copies the Type of Service (TOS) header from the inner IP header to the outer IP header of the encapsulated packets in order to preserve the original TOS information. |
| Step 7 Enable Tunnel Monitoring.<br>You need to assign an IP address to the tunnel interface for monitoring. | To alert the device administrator to tunnel failures and to provide automatic failover to another tunnel interface.<br>1. Specify a **Destination IP** address on the other side of the tunnel to determine if the tunnel is working properly.<br>2. Select a **Profile** to determine the action on tunnel failure. To create a new profile, see Define a Tunnel Monitoring Profile. |
| Step 8 (Required only if the VPN peer uses policy-based VPN). Create a Proxy ID to identify the VPN peers. | 1. Select **Network > IPSec Tunnels > ProxyID**.<br>2. Click **Add** and enter the IP address for the VPN gateway peers.<br> 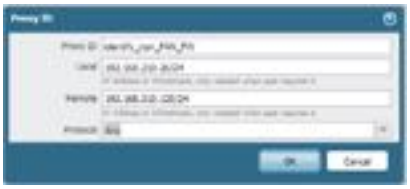 |


© Palo Alto Networks, Inc.

| Set up an IPSec Tunnel | |
|---|---|
| Step 9 Save your changes | Click **OK** and **Commit**. |

# Set up Tunnel Monitoring

To provide uninterrupted VPN service, you can use the Dead Peer Detection capability along with the tunnel monitoring capability on the firewall. You can also monitor the status of the tunnel. For details see:

- Define a Tunnel Monitoring Profile
- View the Status of the Tunnels

## Define a Tunnel Monitoring Profile

A tunnel monitoring profile allows you to verify connectivity between the VPN peers; you can configure the tunnel interface to ping a destination IP address at a specified interval and specify the action if the communication across the tunnel is broken.

| **Define a Tunnel Monitoring Profile** | |
|---|---|
| Step 1 | Select **Network > Network Profiles > Monitor**. A default tunnel monitoring profile is available for use. |
| Step 2 | Click **Add**, and enter a **Name** for the profile. |
| Step 3 | Select the **Action** if the destination IP address is unreachable.<br>• **Wait Recover**—the firewall waits for the tunnel to recover. It continues to use the tunnel interface in routing decisions as if the tunnel is still active.<br>• **Fail Over**—forces traffic to a back-up path if one is available. The firewall disables the tunnel interface, and thereby disables any routes in the routing table that use the interface.<br>In either case, the firewall attempts to accelerate the recovery by negotiating new IPSec keys. |
| Step 4 | Specify the **Interval** and **Threshold** to trigger the specified action.<br>The threshold specifies the number of heartbeats to wait before taking the specified action. The range is 2-100 and the default is 5.<br>The Interval measures the time between heartbeats. The range is 2-10 and the default is 3 seconds. |
| Step 5 | Attach the monitoring profile to the IPsec Tunnel configuration. See Enable Tunnel Monitoring. |

## View the Status of the Tunnels

The status of the tunnel informs you about whether or not valid IKE phase-1 and phase-2 SAs have been established, and whether the tunnel interface is up and available for passing traffic.

Because the tunnel interface is a logical interface, it cannot indicate a physical link status. Therefore, you must enable tunnel monitoring so that the tunnel interface can verify connectivity to an IP address and determine if the path is still usable. If the IP address is unreachable, the firewall will either wait for the tunnel to recover or failover. When a failover occurs, the existing tunnel is torn down and routing changes are triggered to set up a new tunnel and redirect traffic.

| **View Tunnel Status** |
|---|
| 1. Select **Network > IPSec Tunnels**. |

 © Palo Alto Networks, Inc.

**View Tunnel Status**

2. View the **Tunnel Status**.
   - Green indicates a valid IPSec SA tunnel.
   - Red indicates that IPSec SA is not available or has expired.
3. View the **IKE Gateway Status**.
   - Green indicates a valid IKE phase-1 SA.
   - Red indicates that IKE phase-1 SA is not available or has expired.
4. View the **Tunnel Interface Status**.
   - Green indicates that the tunnel interface is up.
   - Red indicates that the tunnel interface is down, because tunnel monitoring is enabled and the status is down.

To troubleshoot a VPN tunnel that is not yet up, see Interpret VPN Error Messages.

# Test VPN Connectivity

## Test Connectivity

- Initiate IKE phase 1 by either pinging a host across the tunnel, or use the following CLI command:

  **test vpn ike-sa gateway** *gateway_name*

- Then enter the following command to test if IKE phase 1 is set up:

  **show vpn ike-sa gateway** *gateway_name*

  In the output check if the Security Association displays. If it does not, review the system log messages to interpret the reason for failure.

- Initiate IKE phase 2 by either pinging a host from across the tunnel, or use the following CLI command:

  **test vpn ipsec-sa tunnel** *tunnel_name*

- Then enter the following command to test if IKE phase 1 is set up:

  **show vpn ipsec-sa tunnel** *tunnel_name*

  In the output check if the Security Association displays. If it does not, review the system log messages to interpret the reason for failure.

- To view the VPN traffic flow information, use the following command:

  **show vpn-flow**

```
admin@PA-500> show vpn flow
total tunnels configured:             1
filter - type IPSec, state any

total IPSec tunnel configured:        1
total IPSec tunnel shown:             1

name               id      state    local-ip       peer-ip       tunnel-i/f
------------------------------------------------------------------------------
vpn-to-siteB       5       active    100.1.1.1     200.1.1.1      tunnel.41
```

 © Palo Alto Networks, Inc.

## Interpret VPN Error Messages

The following table lists some of the common VPN error messages that are logged in the system log.

**Table: Syslog Error Messages for VPN Issues**

| If error is this: | Try this: |
|---|---|
| `IKE phase-1 negotiation is failed as initiator, main mode. Failed SA: x.x.x.x[500]-y.y.y.y[500] cookie:84222f276c2fa2e9:0000000000000000 due to timeout.`<br>or<br>`IKE phase 1 negotiation is failed. Couldn't find configuration for IKE phase-1 request for peer IP x.x.x.x[1929]` | • Verify that the public IP address for each VPN peer is accurate in the IKE Gateway configuration,<br>• Verify that the IP addresses can be pinged and that routing issues are not causing the connection failure. |
| `Received unencrypted notify payload (no proposal chosen) from IP x.x.x.x[500] to y.y.y.y[500], ignored...`<br>or<br>`IKE phase-1 negotiation is failed. Unable to process peer's SA payload.` | Check the IKE Crypto profile configuration to verify that the proposals on both sides have a common encryption, authentication, and DH Group proposal. |
| `pfs group mismatched:my: 2peer: 0`<br>or<br>`IKE phase-2 negotiation failed when processing SA payload. No suitable proposal found in peer's SA payload.` | Check the IPSec Crypto profile configuration to verify that<br>• pfs is either enabled or disabled on both VPN peers<br>• the DH Groups proposed by each peer has at least one DH Group in common |
| `IKE phase-2 negotiation failed when processing Proxy ID. Received local id x.x.x.x/x type IPv4 address protocol 0 port 0, received remote id y.y.y.y/y type IPv4 address protocol 0 port 0.` | The VPN peer on one end is using policy-based VPN. You must configure a Proxy ID on the Palo Alto Networks firewall. See Step 8. |

 © Palo Alto Networks, Inc.

# Site-to-Site VPN Quick Configs

The following sections provide instructions for configuring some common VPN deployments:

- Site-to-Site VPN with Static Routing
- Site-to-Site VPN with OSPF
- Site-to-Site VPN with Static and Dynamic Routing

## Site-to-Site VPN with Static Routing

The following example shows a VPN connection between two sites that use static routes. Without dynamic routing, the tunnel interfaces on VPN Peer A and VPN Peer B do not require an IP address because the firewall automatically uses the tunnel interface as the next hop for routing traffic across the sites. However, to enable tunnel monitoring, a static IP address has been assigned to each tunnel interface.

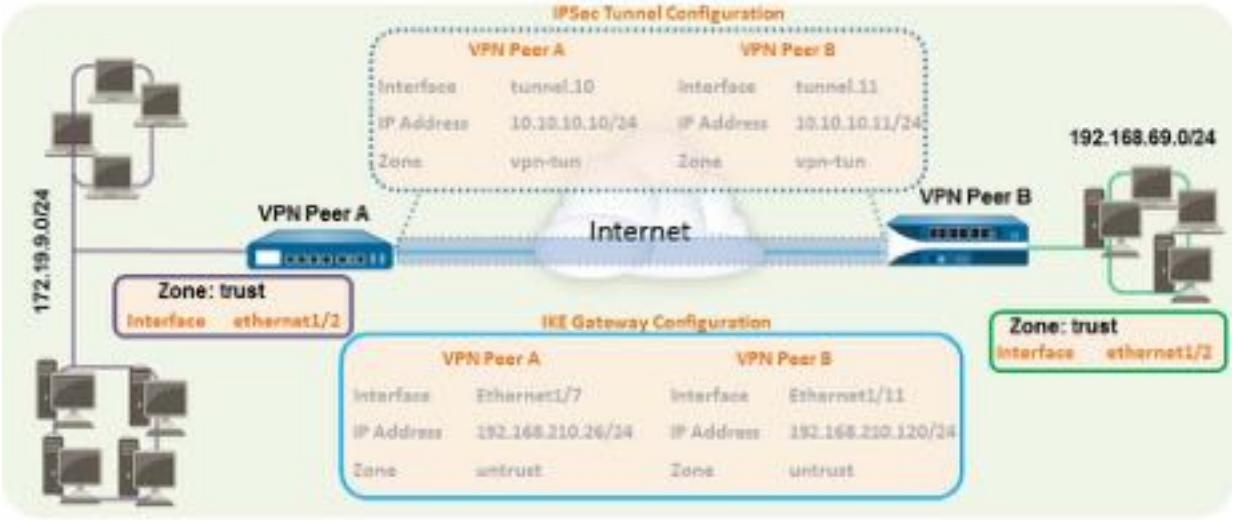


 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static Routing | |
|---|---|
| Step 1 Configure a Layer 3 interface.<br><br>This interface is used for the IKE phase-1 tunnel. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—192.168.210.26/24<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—192.168.210.120/24 |

| Quick Config: Site-to-Site VPN with Static Routing | |
|---|---|
| Step 2 Create a tunnel interface and attach it to a virtual router and security zone.<br> 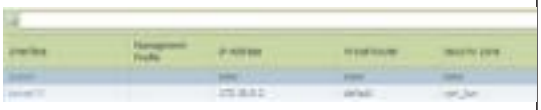 | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, such as **.1**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. (Optional) Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface.<br>With static routes, the tunnel interface does not require an IP address. For traffic that is destined to a specified subnet/IP address, the tunnel interface will automatically become the next hop. Consider adding an IP address if you want to enable tunnel monitoring.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—tunnel.11<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—172.19.9.2/24<br>The configuration for VPN Peer B is:<br>• **Interface**—tunnel.12<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—192.168.69.2/24 |
| Step 3 Configure a static route, on the virtual router, to the destination subnet.<br> | 1. Select **Network > Virtual Router** and click the router you defined in step 4 above.<br>2. Select **Static Route**, click **Add**, and enter a new route to access the subnet that is at the other end of the tunnel.<br>In this example, the configuration for VPN Peer A is:<br>• **Destination**—192.168.69.0/24<br>• **Interface**—tunnel.11<br>The configuration for VPN Peer B is:<br>• **Destination**—172.19.9.0/24<br>• **Interface**—tunnel.12 |


© Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static Routing | | |
|---|---|---|
| Step 4 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).<br><br>Complete this task on both peers and make sure to set identical values. | 1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.<br> <br>2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.<br> |
| Step 5 | Set up the IKE Gateway. | 1. Select **Network > Network Profiles > IKE Gateway.**<br>2. Click **Add** and configure the options in the **General** tab**.**<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Local IP address**—192.168.210.26/24<br>• **Peer IP type/address**—static/192.168.210.120<br>• **Preshared keys**—enter a value<br>• **Local identification**—None; this means that the local IP address will be used as the local identification value.<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Local IP address**—192.168.210.120/24<br>• **Peer IP type/address**—static/192.168.210.26<br>• **Preshared keys**—enter same value as on Peer A<br>• **Local identification**—None<br>3. Select **Advanced Phase 1 Options** and select the IKE Crypto profile you created earlier to use for IKE phase 1. |

| Quick Config: Site-to-Site VPN with Static Routing | |
|---|---|
| Step 6 Set up the IPSec Tunnel.. | 1. Select **Network > IPSec Tunnels**.<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Tunnel Interface**—tunnel.11<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IPSec Crypto profile defined in Step 4.<br>The configuration for VPN Peer B is:<br>• **Tunnel Interface**—tunnel.12<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IPSec Crypto defined in Step 4.<br>3. (Optional) Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity. Typically, the tunnel interface IP address for the VPN Peer is used.<br>4. (Optional) To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile. |
| Step 7 Create policies to allow traffic between the sites (subnets). | 1. Select **Policies > Security**.<br>2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses. |
| Step 8 Save any pending configuration changes. | Click **Commit**. |
| Step 9 Test VPN connectivity. | See View the Status of the Tunnels. |

 © Palo Alto Networks, Inc.

## Site-to-Site VPN with OSPF

In this example, each site uses OSPF for dynamic routing of traffic. The tunnel IP address on each VPN peer is statically assigned and serves as the next hop for routing traffic between the two sites.

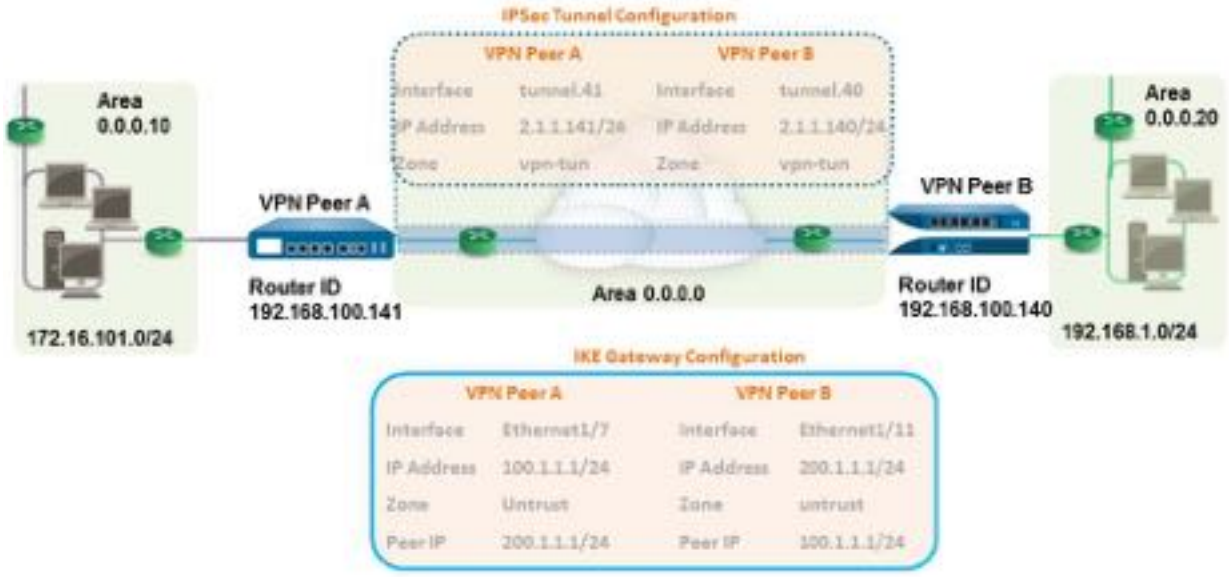

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | |
|---|---|
| Step 1 Configure the Layer 3 interfaces on each firewall. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—100.1.1.1/24<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—200.1.1.1/24 |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | |
|---|---|
| Step 2 Create a tunnel interface and attach it to a virtual router and security zone.<br> | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, say, **.11**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example 172.19.9.2/24.<br>This IP address will be used as the next hop IP address to route traffic to the tunnel and can also be used to monitor the status of the tunnel.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—tunnel.41<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.141/24<br>The configuration for VPN Peer B is:<br>• **Interface**—tunnel.40<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.140/24 |
| Step 3 Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).<br>Complete this task on both peers and make sure to set identical values. | 1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.<br> <br>2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.<br> |

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | | |
|---|---|---|
| Step 4 | Set up the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall.<br><br>For more information on the OSPF options that are available on the firewall, see Configure OSPF.<br><br>Use Broadcast as the link type when there are more than two OSPF routers that need to exchange routing information. | 1. Select **Network > Virtual Routers**, and select the default router or add a new router.<br>2. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.<br>3. In this example, the OSPF configuration for VPN Peer A is:<br>– **Router ID**: 192.168.100.141<br>– **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p<br>– **Area ID**: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast<br>The OSPF configuration for VPN Peer B is:<br>– **Router ID**: 192.168.100.140<br>– **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p<br>– **Area ID**: 0.0.0.20 that is assigned to the interface Ethernet1/15 and Link Type: Broadcast |
| Step 5 | Set up the IKE Gateway.<br><br>This examples uses static IP addresses for both VPN peers. Typically, the corporate office uses a statically configured IP address, and the branch side can be a dynamic IP address; dynamic IP address are not best suited for configuring stable services such as VPN. | 1. Select **Network > Network Profiles > IKE Gateway.**<br>2. Click **Add** and configure the options in the **General** tab**.**<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Local IP address**—100.1.1.1/24<br>• **Peer IP address**—200.1.1.1/24<br>• **Preshared keys**—enter a value<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Local IP address**—200.1.1.1/24<br>• **Peer IP address**—100.1.1.1/24<br>• **Preshared keys**—enter same value as on Peer A<br>3. Select the IKE Crypto profile you created earlier to use for IKE phase 1. |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | | |
|---|---|---|
| Step 6 | Set up the IPSec Tunnel. | 1. Select **Network > IPSec Tunnels**.<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Tunnel Interface**—tunnel.41<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>The configuration for VPN Peer B is:<br>• **Tunnel Interface**—tunnel.40<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.<br>4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile. |
| Step 7 | Create policies to allow traffic between the sites (subnets). | 1. Select **Policies > Security**.<br>2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses. |

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | |
|---|---|
| Step 8 Verify OSPF adjacencies and routes from the CLI. | Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.<br>• **show routing protocol ospf neighbor**<br> 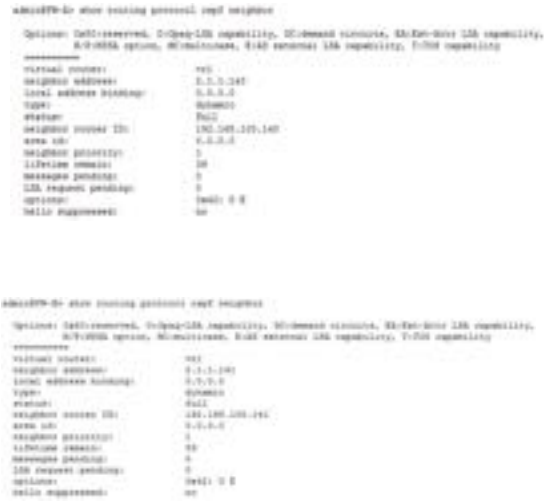 |
| | • **show routing route type ospf**<br> 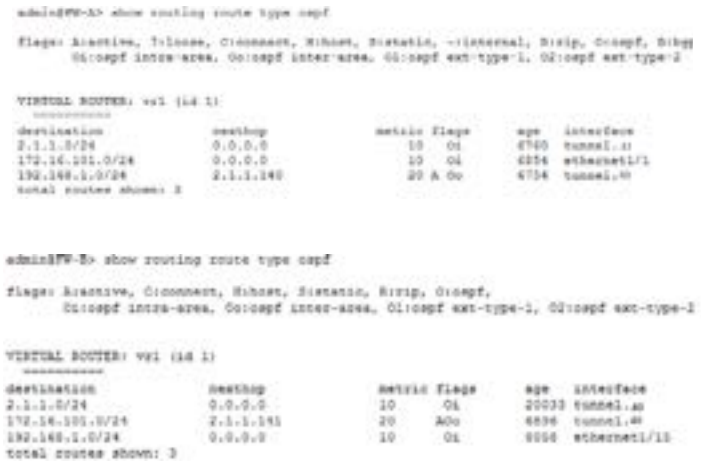 |
| Step 9 Test VPN connectivity. | See Set up Tunnel Monitoring and View the Status of the Tunnels. |


© Palo Alto Networks, Inc.

## Site-to-Site VPN with Static and Dynamic Routing

In this example, one site uses static routes and the other site uses OSPF. When the routing protocol is not the same between the locations, the tunnel interface on each firewall must be configured with a static IP address. Then, to allow the exchange of routing information, the firewall that participates in both the static and dynamic routing process must be configured with a *Redistribution profile*. Configuring the redistribution profile, enables the virtual router to redistribute and filter routes between protocols—static routes, connected routes, and hosts—from the static autonomous system to the OSPF autonomous system. Without this redistribution profile, each protocol functions on its own and does not exchange any route information with other protocols running on the same virtual router.

In this example, the satellite office has static routes and all traffic destined to the 192.168.x.x network is routed to tunnel.41. The virtual router on VPN Peer B participates in both the static and the dynamic routing process and is configured with a redistribution profile in order to propagate (*export*) the static routes to the OSPF autonomous system.

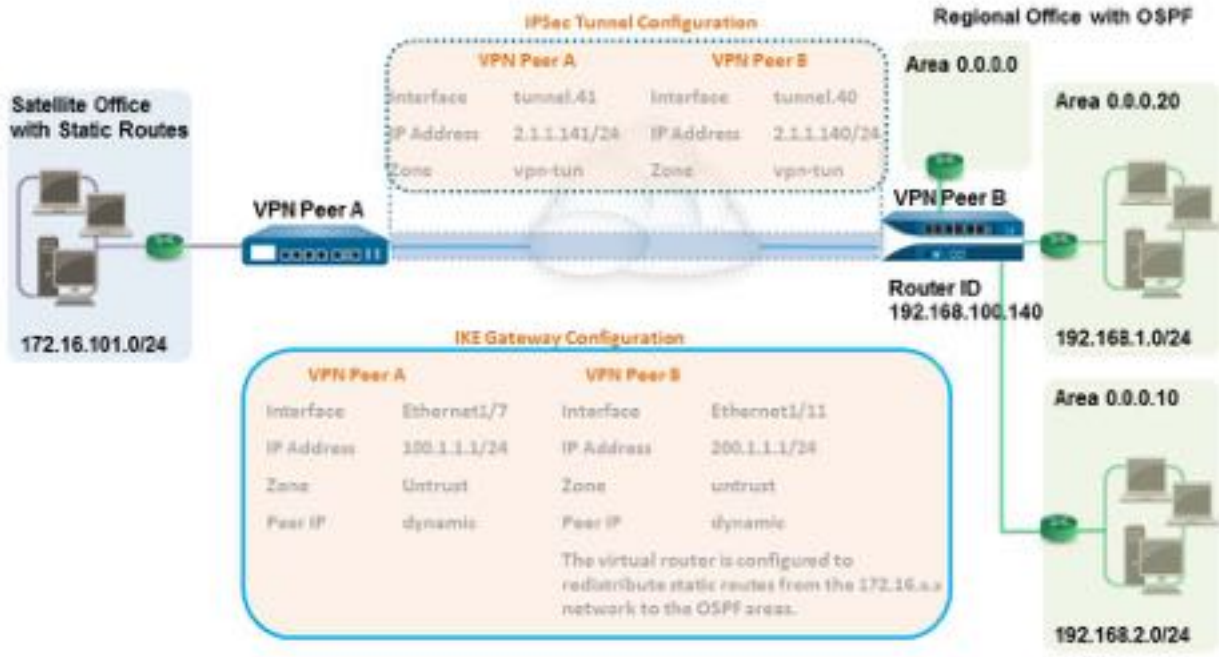

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | | |
|---|---|---|
| Step 1 | Configure the Layer 3 interfaces on each firewall. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—100.1.1.1/24<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—200.1.1.1/24 |
| Step 2 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).<br>Complete this task on both peers and make sure to set identical values. | 1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.<br> <br>2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.<br> |


© Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 3 Set up the IKE Gateway.<br><br>With pre-shared keys, to add authentication scrutiny when setting up the IKE phase-1 tunnel, you can set up Local and Peer Identification attributes and a corresponding value that is matched in the IKE negotiation process. | 1. Select **Network > Network Profiles > IKE Gateway.**<br>2. Click **Add** and configure the options in the **General** tab**.**<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Local IP address**—100.1.1.1/24<br>• **Peer IP type**—dynamic<br>• **Preshared keys**—enter a value<br>• **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer A.<br>• **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer B<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Local IP address**—200.1.1.1/24<br>• **Peer IP address**—dynamic<br>• **Preshared keys**—enter same value as on Peer A<br>• **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer B<br>• **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer A<br>3. Select the IKE Crypto profile you created earlier to use for IKE phase 1. |

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 4 Create a tunnel interface and attach it to a virtual router and security zone.<br> | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, say, **.41**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example 172.19.9.2/24.<br>This IP address will be used to route traffic to the tunnel and to monitor the status of the tunnel.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—tunnel.41<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.141/24<br>The configuration for VPN Peer B is:<br>• **Interface**—tunnel.42<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.140/24 |
| Step 5 Specify the interface to route traffic to a destination on the 192.168.x.x network. | 1. On VPN Peer A, select the virtual router.<br>2. Select **Static Routes**, and **Add** tunnel.41 as the **Interface** for routing traffic with a **Destination** in the 192.168.x.x network. |


© Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 6 Set up the static route and the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall. | 1. On VPN Peer B, select **Network > Virtual Routers**, and select the default router or add a new router.<br>2. Select **Static Routes** and **Add** the tunnel IP address as the next hop for traffic in the 172.168.x.x. network.<br>Assign the desired route metric; using a lower the value makes the a higher priority for route selection in the forwarding table.<br>3. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.<br>4. In this example, the OSPF configuration for VPN Peer B is:<br>• Router ID: 192.168.100.140<br>• Area ID: 0.0.0.0 is assigned to the interface Ethernet 1/12 Link type: Broadcast<br>• Area ID: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast<br>• Area ID: 0.0.0.20 is assigned to the interface Ethernet1/15 and Link Type: Broadcast |
| Step 7 Create a redistribution profile to inject the static routes into the OSPF autonomous system.<br> <br> | 1. Create a redistribution profile on VPN Peer B.<br>a. Select **Network > Virtual Routers**, and select the router you used above.<br>b. Select **Redistribution Profiles,** and click **Add**.<br>c. Enter a Name for the profile and select **Redist** and assign a **Priority** value. If you have configured multiple profiles, the profile with the lowest priority value is matched first.<br>d. Set **Source Type** as **static**, and click **OK**. The static route defined in Step 6-2 will be used for the redistribution.<br>2. Inject the static routes in to the OSPF system.<br>a. Select **OSPF> Export Rules** (for IPv4) or **OSPFv3> Export Rules** (for IPv6).<br>b. Click **Add**, and select the redistribution profile that you just created.<br>c. Select how the external routes are brought into the OSPF system. The default option, **Ext2** calculates the total cost of the route using only the external metrics. To use both internal and external OSPF metrics, use **Ext1.**<br>d. Assign a **Metric** (cost value) for the routes injected into the OSPF system. This option allows you to change the metric for the injected route as it comes into the OSPF system.<br>e. Click **OK** to save the changes. |

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 8 Set up the IPSec Tunnel. | 1. Select **Network > IPSec Tunnels**.<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Tunnel Interface**—tunnel.41<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>The configuration for VPN Peer B is:<br>• **Tunnel Interface**—tunnel.40<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.<br>4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile. |
| Step 9 Create policies to allow traffic between the sites (subnets). | 1. Select **Policies > Security**.<br>2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses. |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 10 Verify OSPF adjacencies and routes from the CLI. | Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.<br>• **`show routing protocol ospf neighbor`**<br> 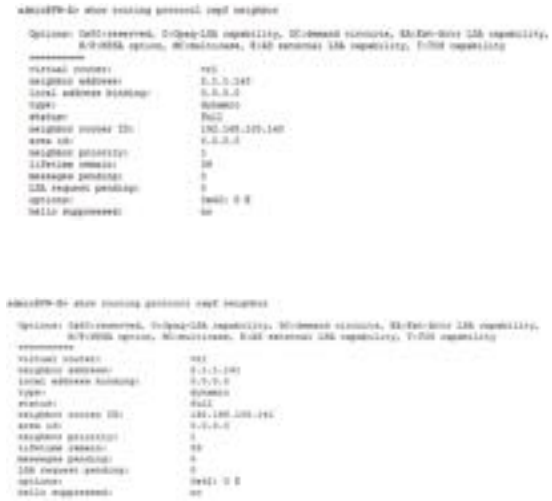 |
| | • **show routing route**<br>The following is an example of the output on each VPN peer.<br> |
| Step 11 Test VPN connectivity. | See Set up Tunnel Monitoring and View the Status of the Tunnels. |

 © Palo Alto Networks, Inc.

# Large Scale VPN (LSVPN)

The GlobalProtect Large Scale VPN (LSVPN) feature on the Palo Alto Networks next-generation firewall simplifies the deployment of traditional hub and spoke VPNs, enabling you to quickly deploy enterprise networks with several branch offices with a minimum amount of configuration required on the remote *satellite* devices. This solution uses certificates for device authentication and IPSec to secure data.

LSVPN enables site-to-site VPNs between Palo Alto Networks firewalls. To set up a site-to-site VPN between a Palo Alto Networks firewall and another device, see VPNs.

The following topics describe the LSVPN components and how to set them up to enable site-to-site VPN services between Palo Alto Networks firewalls:

- LSVPN Overview
- Create Interfaces and Zones for the LSVPN
- Enable SSL Between GlobalProtect LSVPN Components
- Configure the Portal to Authenticate Satellites
- Configure GlobalProtect Gateways for LSVPN
- Configure the GlobalProtect Portal for LSVPN
- Prepare the Satellite Device to Join the LSVPN
- Verify the LSVPN Configuration
- LSVPN Quick Configs

# LSVPN Overview

GlobalProtect provides a complete infrastructure for managing secure access to corporate resources from your remote sites. This infrastructure includes the following components:

- **GlobalProtect Portal**—Provides the management functions for your GlobalProtect LSVPN infrastructure. Every satellite that participates in the GlobalProtect LSVPN receives configuration information from the portal, including configuration information to enable the satellites (the spokes) to connect to the gateways (the hubs). You configure the portal on an interface on any Palo Alto Networks next-generation firewall.
- **GlobalProtect Gateways**—A Palo Alto Networks firewall that provides the tunnel end point for satellite connections. The resources that the satellites access is protected by security policy on the gateway. It is not required to have a separate portal and gateway; a single firewall can function both as portal and gateway.
- **GlobalProtect Satellite**—A Palo Alto Networks firewall at a remote site that establishes IPSec tunnels with the gateway(s) at your corporate office(s) for secure access to centralized resources. Configuration on the satellite firewall is minimal, enabling you to quickly and easily scale your VPN as you add new sites.

The following diagram illustrates how the GlobalProtect LSVPN components work together.

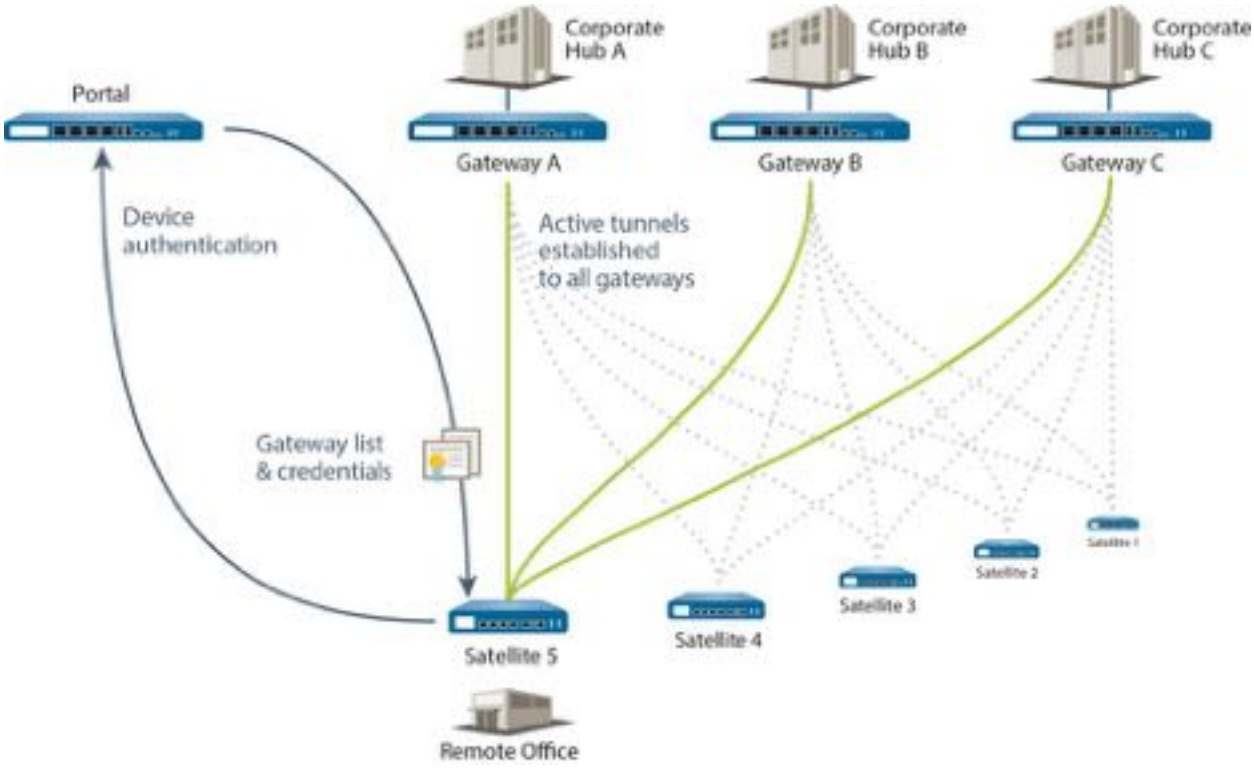


 © Palo Alto Networks, Inc.

# Create Interfaces and Zones for the LSVPN

You must configure the following interfaces and zones for your LSVPN infrastructure:

- **GlobalProtect portal**—Requires a Layer 3 interface for GlobalProtect satellites to connect to. If the portal and gateway are on the same firewall, they can use the same interface. The portal must be in a zone that is accessible from your branch offices.
- **GlobalProtect gateways**—Requires three interfaces: a Layer 3 interface in the zone that is reachable by the remote satellites, an internal interface in the trust zone that connects to the protected resources, and a logical tunnel interface for terminating the VPN tunnels from the satellites. Unlike other site-to-site VPN solutions, the GlobalProtect gateway only requires a single tunnel interface, which it will use for tunnel connections with all of your remote satellites (point-to-multipoint). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface.
- **GlobalProtect satellites**—Requires a single tunnel interface for establishing a VPN with the remote gateways (up to a maximum of 25 gateways). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface.

For more information about portals, gateways, and satellites see LSVPN Overview.

| Set Up Interfaces and Zones for the GlobalProtect LSVPN | |
|---|---|
| Step 1 Configure a Layer 3 interface.<br><br>The portal and each gateway and satellite all require a Layer 3 interface to enable traffic to be routed between sites.<br><br>If the gateway and portal are on the same firewall, you can use a single interface for both components.<br><br>IPv6 addresses are not supported with LSVPN. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for GlobalProtect LSVPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.11.100/24.<br>6. To save the interface configuration, click **OK**. |

**Set Up Interfaces and Zones for the GlobalProtect LSVPN (Continued)**

| | | |
|---|---|---|
| Step 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.<br><br>IP addresses are not required on the tunnel interface unless plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.<br><br>Make sure to enable User-ID in the zone where the VPN tunnels terminate. | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *lsvpn-tun*), select the **Enable User Identification** check box, and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. (Optional) If you want to assign an IP address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.11.33/24.<br>6. To save the interface configuration, click **OK**. |
| Step 3 | If you created a separate zone for tunnel termination of VPN connections, create a security policy to enable traffic flow between the VPN zone and your trust zone. For example, the following policy rule enables traffic between the *lsvpn-tun* zone and the *L3-Trust* zone.<br> 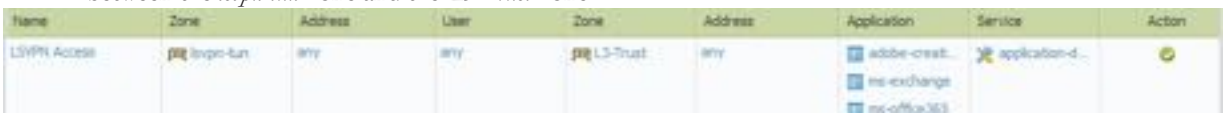 <br> | |
| Step 4 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# Enable SSL Between GlobalProtect LSVPN Components

All interaction between the GlobalProtect components occurs over an SSL connection. Therefore, you must generate and/or install the required certificates before configuring each component so that you can reference the appropriate certificate(s) and/or certificate profiles in the configurations for each component. The following sections describe the supported methods of certificate deployment, descriptions and best practice guidelines for the various GlobalProtect certificates, and provide instructions for generating and deploying the required certificates:

▲ About Certificate Deployment

▲ Deploy Server Certificates to the GlobalProtect LSVPN Components

## About Certificate Deployment

There are two basic approaches to deploying certificates for GlobalProtect LSVPN:

- **Enterprise Certificate Authority**—If you already have your own enterprise certificate authority, you can use this internal CA to issue an intermediate CA certificate for the GlobalProtect portal to enable it to issue certificates to the GlobalProtect gateways and satellites.
- **Self-Signed Certificates**—You can generate a self-signed root CA certificate on the firewall and use it to issue server certificates for the portal, gateway(s), and satellite(s). As a best practice, create a self-signed root CA certificate on the portal and use it to issue server certificates for the gateways and satellites. This way, the private key used for certificate signing stays on the portal.

## Deploy Server Certificates to the GlobalProtect LSVPN Components

The GlobalProtect LSVPN components use SSL/TLS to mutually authenticate. Before deploying the LSVPN, you must issue server certificates to the portal and gateways. You do not need to create server certificates for the satellite devices because the portal will issue a server certificate for each satellite during the first connection.

In addition, you must import the root CA certificate used to issue the server certificates onto each firewall that you plan to host as a gateway or satellite. Finally, on each gateway and satellite participating in the LSVPN, you must configure a certificate profile that will enable them to establish an SSL/TLS connection using mutual authentication.

The following workflow shows the best practice steps for deploying SSL certificates to the GlobalProtect LSVPN components:

You do not need to issue server certificates for the satellite devices because the portal will issue them as part of the satellite registration process.

| Deploy SSL Server Certificates to the GlobalProtect Components | | |
|---|---|---|
| Step 1 | On the firewall hosting the portal, create the root CA certificate for issuing self-signed certificates for the GlobalProtect components. | To use self-signed certificates, you must first create the root CA certificate that will be used to sign the certificates for the GlobalProtect components as follows:<br>1. To create a root CA certificate, select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**.<br>2. Enter a **Certificate Name**, such as *LSVPN_CA*. The certificate name cannot contain any spaces.<br>3. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed).<br>4. Select the **Certificate Authority** check box and then click **OK** to generate the certificate. |
| Step 2 | Generate server certificates for the GlobalProtect portal and gateway(s).<br>You must issue a unique self-signed server certificate for the portal and for each GlobalProtect gateway. The best practice is to issue all of the required certificates on the portal, so that the signing certificate (with the private key) does not have to be exported.<br> If the GlobalProtect portal and gateway are on the same firewall interface, you can use the same server certificate for both components. | Use the root CA on the portal to generate server certificates for each gateway you plan to deploy:<br>1. Select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**.<br>2. Enter a **Certificate Name**. The Certificate Name cannot contain any spaces.<br>3. Enter the FQDN (recommended) or IP address of the interface where you plan to configure the gateway in the **Common Name** field.<br>4. In the **Signed By** field, select the *LSVPN_CA* you created previously.<br>5. In the Certificate Attributes section, click **Add** and define the attributes to uniquely identify the gateway. Keep in mind that if you add a **Host Name** attribute (which populates the SAN field of the certificate), it must exactly match the value you defined for the **Common Name**.<br>6. **Generate** the certificate. |

 © Palo Alto Networks, Inc.

| Deploy SSL Server Certificates to the GlobalProtect Components (Continued) | |
|---|---|
| Step 3 Deploy the self-signed server certificates to the gateways.<br><br>**Best Practices:**<br>• Export the self-signed server certificates issued by the root CA from the portal and import them onto the gateways.<br>• Be sure to issue a unique server certificate for each gateway.<br>• The Common Name (CN) and, if applicable, the Subject Alternative Name (SAN) fields of the certificate must match the IP address or fully qualified domain name (FQDN) of the interface where you configure the gateway. | 1. On the portal, select **Device > Certificate Management > Certificates > Device Certificates**, select the gateway certificate you want to deploy, and click **Export**.<br>2. Select **Encrypted Private Key and Certificate (PKCS12)** from the **File Format** drop-down.<br>3. Enter (and re-enter) a **Passphrase** to encrypt the private key associated with the certificate and then click **OK** to download the PKCS12 file to your computer.<br>4. On the gateway, select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>5. Enter a **Certificate Name**.<br>6. Enter the path and name to the **Certificate File** you just downloaded from the portal, or **Browse** to find the file.<br>7. Select **Encrypted Private Key and Certificate (PKCS12)** as the **File Format**.<br>8. Enter the path and name to the PKCS12 file in the **Key File** field or **Browse** to find it.<br>9. Enter and re-enter the **Passphrase** you used to encrypt the private key when you exported it from the portal and then click **OK** to import the certificate and key. |
| Step 4 Import the root CA certificate used to issue server certificates for the LSVPN components.<br><br>You must import the root CA certificate onto all gateways and satellites. For security reasons, make sure you export the certificate only, and not the associated private key. | 1. Download the root CA certificate from the portal.<br>a. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>b. Select the root CA certificate used to issue certificates for the LSVPN components and click **Export**.<br>c. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (Do not export the private key.)<br>2. On the firewalls hosting the gateways and satellites, import the root CA certificate.<br>a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.<br>c. **Browse** to the **Certificate File** you downloaded from the CA.<br>d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.<br>e. Select the certificate you just imported on the **Device Certificates** tab to open it.<br>f. Select **Trusted Root CA** and then click **OK**.<br>g. **Commit** the changes. |

| Deploy SSL Server Certificates to the GlobalProtect Components (Continued) | | |
|---|---|---|
| Step 5 | Create a certificate profile.<br>The GlobalProtect LSVPN portal and each gateway require a certificate profile that specifies which certificate to use to authenticate the satellites. | 1. Select **Device > Certificate Management > Certificate Profile** and click **Add** and enter a profile **Name**.<br>2. Make sure **Username Field** is set to **None**.<br>3. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you imported in Step 4.<br>4. (Optional, but recommended) Enable use of CRL and/or OCSP to enable certificate status verification.<br>5. Click **OK** to save the profile. |
| Step 6 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

# Configure the Portal to Authenticate Satellites

In order to register with the LSVPN, each satellite must establish an SSL/TLS connection with the portal. After establishing the connection, the portal authenticates the satellite device to ensure that is authorized to join the LSVPN. After successfully authenticating the satellite, the portal will issue a server certificate for the satellite and push the LSVPN configuration specifying the gateways to which the satellite can connect and the root CA certificate required to establish an SSL connection with the gateways.

There are two ways that the satellite can authenticate to the portal during its initial connection:

- **Serial number**—You can configure the portal with the serial number of the satellite firewalls that are authorized to join the LSVPN. During the initial satellite connection to the portal, the satellite presents its serial number to the portal and if the portal has the serial number in its configuration, the satellite will be successfully authenticated. You add the serial numbers of authorized satellites when you configure the portal. See Configure the Portal.
- **Username and password**—If you would rather provision your satellites without manually entering the serial numbers of the satellite devices into the portal configuration, you can instead require the satellite administrator to authenticate when establishing the initial connection to the portal. Although the portal will always look for the serial number in the initial request from the satellite, if it cannot identify the serial number, the satellite administrator must provide a username and password to authenticate to the portal. Because the portal will always fall back to this form of authentication, you must create an authentication profile in order to commit the portal configuration. This requires that you set up an authentication profile for the portal LSVPN configuration even if you plan to authenticate satellites using the serial number.

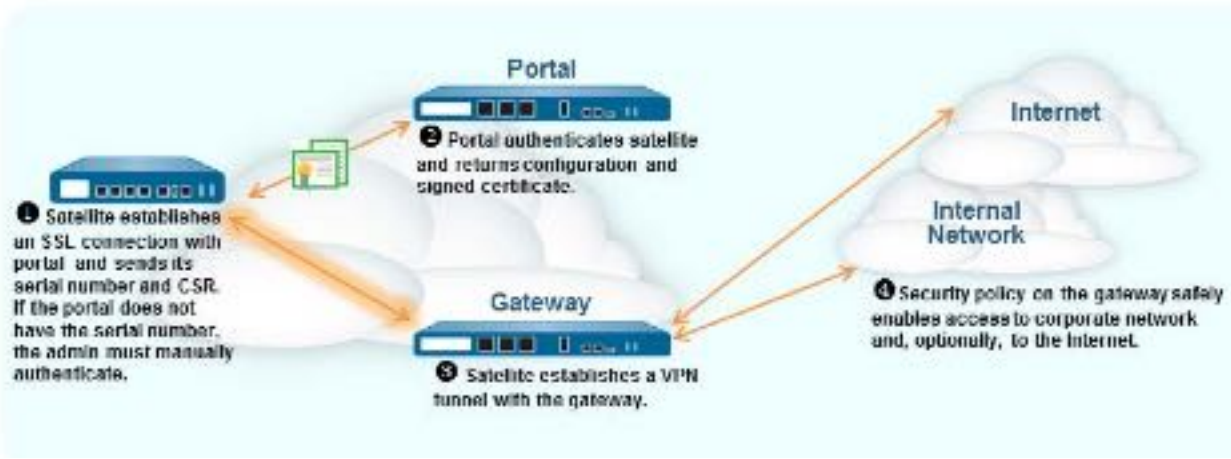


The following workflow describes how to set up the portal to authenticate satellites against an existing authentication service. GlobalProtect LSVPN supports external authentication using a local database, LDAP (including Active Directory), Kerberos, or RADIUS.

| Set Up Satellite Authentication | | |
|---|---|---|
| Step 1 | Create a server profile on the portal.<br><br>The server profile instructs the firewall how to connect to an external authentication service to validate the authentication credentials provided by the satellite administrator.<br><br>If you are using local authentication, you can skip this step and instead add a local user configuration for authenticating the satellite administrator.<br><br> If you are using LDAP to connect to Active Directory (AD), you must create a separate LDAP server profile for every AD domain. | 1. Select **Device > Server Profiles** and select type of profile (**LDAP**, **Kerberos**, or **RADIUS**).<br>2. Click **Add** and enter a **Name** for the profile, such as *LSVPN-Auth*.<br>3. (LDAP only) Select the **Type** of LDAP server you are connecting to.<br>4. Click **Add** in the Servers section and then enter information required to connect to the authentication service, including the server **Name**, **IP Address** (or FQDN), and **Port**.<br>5. (RADIUS and LDAP only) Specify settings to enable the firewall to authenticate to the authentication service as follows:<br>• RADIUS—Enter the shared **Secret** when adding the server entry.<br>• LDAP—Enter the **Bind DN** and **Bind Password**.<br>6. (LDAP and Kerberos only) Specify where to search for credentials in the directory service:<br>• LDAP—The **Base** DN specifies where in the LDAP tree to begin searching for users and groups. This field should populate automatically when you enter the server address and port. If it doesn't, check the service route to the LDAP server.<br>• Kerberos—Enter the Kerberos **Realm** name.<br>7. Specify the **Domain** name (without dots, for example `acme` not `acme.com`).<br>8. Click **OK** to save the server profile. |
| Step 2 | Create an authentication profile.<br><br>The authentication profile specifies which server profile to use to authenticate satellites. | 1. Select **Device > Authentication Profile** and click **Add**.<br>2. Enter a **Name** for the profile and then select the **Authentication** type (**Local Database, LDAP**, **Kerberos**, or **RADIUS**).<br>3. Select the **Server Profile** you created in Step 1.<br>4. (LDAP AD) Enter **sAMAccountName** as the **Login Attribute**.<br>5. Click **OK**. |
| Step 3 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# Configure GlobalProtect Gateways for LSVPN

Because the GlobalProtect configuration that the portal delivers to the satellites includes the list of gateways the satellite can connect to, it is a good idea to configure the gateways before configuring the portal.


- ▲ Prerequisite Tasks
- ▲ Configure the Gateway


## Prerequisite Tasks

Before you can configure the GlobalProtect gateway, you must have completed the following tasks:

- ❒ Created the interfaces (and zones) for the interface where you plan to configure each gateway. You must configure both the physical interface and the virtual tunnel interface. See Create Interfaces and Zones for the LSVPN.
- ❒ Set up the gateway server certificates and certificate profile required for enable GlobalProtect satellite and gateway to establish a mutual SSL/TLS connection. See Enable SSL Between GlobalProtect LSVPN Components.

## Configure the Gateway

After you have completed the Prerequisite Tasks, configure each GlobalProtect gateway to participate in the LSVPN as follows:

| Configure the Gateway for LSVPN | | |
|---|---|---|
| Step 1 | Add a gateway. | 1. Select **Network > GlobalProtect > Gateways** and click **Add**.<br>2. On the **General** tab, enter a **Name** for the gateway. The gateway name should not contain any spaces and as a best practice it should include the location or other descriptive information that will help identify the gateway.<br>3. (Optional) Select the virtual system to which this gateway belongs from the **Location** field. |
| Step 2 | Specify the network information to enable satellites to connect to the gateway.<br><br>If you have not yet created the network interface for the gateway, see Create Interfaces and Zones for the LSVPN for instructions. If you haven't yet created a server certificate for the gateway, see Deploy Server Certificates to the GlobalProtect LSVPN Components. | 1. Select the **Interface** that satellites will use for ingress access to the gateway.<br>2. Select the **IP Address** for gateway access.<br>3. Select the **Server Certificate** for the gateway from the drop-down. |

| Configure the Gateway for LSVPN (Continued) | | |
|---|---|---|
| Step 3 | Select the certificate profile for the gateway to use to authenticate satellites attempting to establish tunnels.<br><br>If you have not yet set up the certificate profile, see Enable SSL Between GlobalProtect LSVPN Components for instructions. | Select the **Certificate Profile** to you created for SSL communication between the LSVPN components. |
| Step 4 | Configure the tunnel parameters and enable tunneling. | 1. On the GlobalProtect Gateway dialog, select **Satellite Configuration > Tunnel Settings**.<br>2. Select the **Tunnel Configuration** check box to enable tunneling.<br>3. Select the **Tunnel Interface** you defined in Step 2 in Create Interfaces and Zones for the LSVPN.<br>4. (Optional) If you want to preserve the Type of Service (ToS) information in the encapsulated packets, select the **Copy TOS** check box. |
| Step 5 | (Optional) Enable tunnel monitoring.<br><br>Tunnel monitoring enables satellite devices to monitor its gateway tunnel connection, allowing it to failover to a backup gateway if the connection fails. Failover to another gateway is the only type of tunnel monitoring profile supported with LSVPN. | 1. Select the **Tunnel Monitoring** check box.<br>2. Specify the **Destination IP** address the satellite devices should use to determine if the gateway is active. Alternatively, if you configured an IP address for the tunnel interface, you can leave this field blank and the tunnel monitor will instead use the tunnel interface to determine if the connection is active.<br>3. Select **Failover** from the **Tunnel Monitor Profile** drop-down (this is the only supported tunnel monitor profile for LSVPN). |
| Step 6 | Select the Crypto Profile to use when establishing tunnel connections.<br><br>The crypto profile specifies the type IPSec encryption and/or the authentication method for securing the data that will traverse the tunnel. Because both tunnel endpoints in an LSVPN are trusted firewalls within your organization, you can typically use the default profile, which uses ESP-DH group2-AES 128 with SHA-1 encryption.<br><br>However, if you require a different mix of encryption and authentication mechanisms, you can optionally create a custom IPSec crypto profile. | Select **default** from the **IPSec Crypto Profile** drop-down or, optionally, select **New IPSec Crypto Profile** to define a new profile. For details on the authentication and encryption options in the crypto profile, refer to the online help. |

 © Palo Alto Networks, Inc.

| Configure the Gateway for LSVPN (Continued) | | |
|---|---|---|
| Step 7 | Configure the network settings to assign the satellites during establishment of the IPSec tunnel.<br><br> You can also configure the satellite device to push the DNS settings to its local clients by configuring a DHCP server on the firewall hosting the satellite. In this configuration, the satellite will push DNS settings it learns from the gateway to the DHCP clients. | 1. On the GlobalProtect Gateway dialog, select **Satellite Configuration > Network Settings**.<br>2. (Optional) If clients local to the satellite device need to resolve FQDNs on the corporate network, configure the gateway to push DNS settings to the satellites in one of the following ways:<br>• Manually define the **Primary DNS**, **Secondary DNS**, and **DNS Suffix** settings to push to the satellites.<br>• If the gateway has an interface that is configured as a DHCP client, you can set the **Inheritance Source** to that interface and the GlobalProtect satellites will be assigned the same settings received by the DHCP client.<br>3. To specify the **IP Pool** of addresses to assign the tunnel interface on the satellite devices when the VPN is established, click **Add** and then specify the IP address range(s) to use. If you are using dynamic routing, make sure that the IP address pool you designate for satellites does not overlap with the IP addresses you manually assigned to the tunnel interfaces on your gateways and satellites.<br>4. To define what destination subnets to route through the tunnel click **Add** in the **Access Route** area and then enter the routes as follows:<br>• If you want to route all traffic from the satellites through the tunnel, leave this field blank. Note that in this case, all traffic except traffic destined for the local subnet will be tunneled to the gateway.<br>• To route only some traffic through the gateway (called *split tunneling*), specify the destination subnets that must be tunneled. In this case, the satellite will route traffic that is not destined for a specified access route using its own routing table. For example, you may choose to only tunnel traffic destined for your corporate network, and use the local satellite to safely enable Internet access.<br>• If you want to enable routing between satellites, enter the summary route for the network protected by each satellite. |
| Step 8 | (Optional) Define what routes, if any, the gateway will accept from satellites.<br><br>By default, the gateway will not add any routes satellites advertise to its routing table. If you do not want the gateway to accept routes from gateways, you do not need to complete this step. | 1. To enable the gateway to accept routes advertised by satellites, select **Satellite Configuration > Route Filter**.<br>2. Select the **Accept published routes** check box.<br>3. To filter which of the routes advertised by the satellites to add to the gateway routing table, click **Add** and then define the subnets to include. For example, if all the satellites are configured with subnet 192.168.x.0/24 on the LAN side, configuring a permitted route of 192.168.0.0/16 to enable the gateway to only accept routes from the satellite if it is in the 192.168.0.0/16 subnet. |

| Configure the Gateway for LSVPN (Continued) | |
|---|---|
| Step 9 Save the gateway configuration. | 1. Click **OK** to save the settings and close the GlobalProtect Gateway dialog.<br>2. **Commit** the configuration. |

 © Palo Alto Networks, Inc.

# Configure the GlobalProtect Portal for LSVPN

The GlobalProtect portal provides the management functions for your GlobalProtect LSVPN. Every satellite system that participates in the LSVPN receives configuration information from the portal, including information about available gateways as well as the certificate it needs in order to connect to the gateways.

The following sections provide procedures for setting up the portal:

- Prerequisite Tasks
- Configure the Portal
- Define the Satellite Configurations

## Prerequisite Tasks

Before you can configure the GlobalProtect portal, you must have completed the following tasks:

- ☐ Created the interfaces (and zones) for the firewall interface where you plan to configure the portal. See Create Interfaces and Zones for the LSVPN.
- ☐ Issued the portal server certificate, gateway server certificates, and set up the portal to issue server certificates for the satellites. See Enable SSL Between GlobalProtect LSVPN Components.
- ☐ Defined the authentication profile that will be used to authenticate GlobalProtect satellites in the event that the serial number is not available. See Configure the Portal to Authenticate Satellites.
- ☐ Configured the global protect gateways. See Configure GlobalProtect Gateways for LSVPN.

## Configure the Portal

After you have completed the Prerequisite Tasks, configure the GlobalProtect portal as follows:

| **Configure the Portal for LSVPN** | |
|---|---|
| Step 1 Add the portal. | 1. Select **Network > GlobalProtect > Portals** and click **Add**.<br>2. On the **Portal Configuration** tab, enter a **Name** for the portal. The portal name should not contain any spaces.<br>3. (Optional) Select the virtual system to which this portal belongs from the **Location** field. |

| Configure the Portal for LSVPN (Continued) | | |
|---|---|---|
| Step 2 | Specify the network information to enable satellites to connect to the portal.<br><br>If you have not yet created the network interface for the portal, see Create Interfaces and Zones for the LSVPN for instructions. If you haven't yet created a server certificate for the portal and issued gateway certificates, see Deploy Server Certificates to the GlobalProtect LSVPN Components. | 1. Select the **Interface** that satellites will use for ingress access to the portal.<br>2. Select the **IP Address** for satellite access to the portal.<br>3. Select the **Server Certificate** you generated to enable the satellite to establish SSL connection with the portal. |
| Step 3 | Specify an authentication profile for authenticating satellite devices.<br><br>Even if you plan to manually configure the portal with the serial numbers of the satellites, you must define an authentication profile or you will not be able to save the configuration. | • Select the **Authentication Profile** you defined for authenticating satellites.<br>• If you have not yet set up the authentication profile, select **New Authentication Profile** to create one now. See Configure the Portal to Authenticate Satellites for instructions. If the portal is unable to validate the serial number of a connecting satellite, it will fall back to the authentication profile and therefore you must configure an authentication profile in order to save the portal configuration. |
| Step 4 | Continue with defining the configurations to push to the satellites, or, if you have already created the satellite configurations, save the portal configuration. | Click **OK** to save the portal configuration or continue to Define the Satellite Configurations. |

## Define the Satellite Configurations

When a GlobalProtect satellite connects and successfully authenticates to the GlobalProtect portal, the portal delivers a satellite configuration, which specifies what gateways the satellite can connect to. If all your satellites will use the same gateway and certificate configurations, you can create a single satellite configuration to deliver to all satellites upon successful authentication. However, if you require different satellite configurations—for example if you want one group of satellites to connect to one gateway and another group of satellites to connect to a different gateway—you can create a separate satellite configuration for each. The portal will then use the enrollment username/group name or the serial number of the satellite device to determine which satellite configuration to deploy. As with security rule evaluation, the portal looks for a match starting from the top of the list. When it finds a match, it delivers the corresponding configuration to the satellite.

For example, the following figure shows a network in which some branch offices require VPN access to the corporate applications protected by your perimeter firewalls and another site needs VPN access to the datacenter.

 © Palo Alto Networks, Inc.

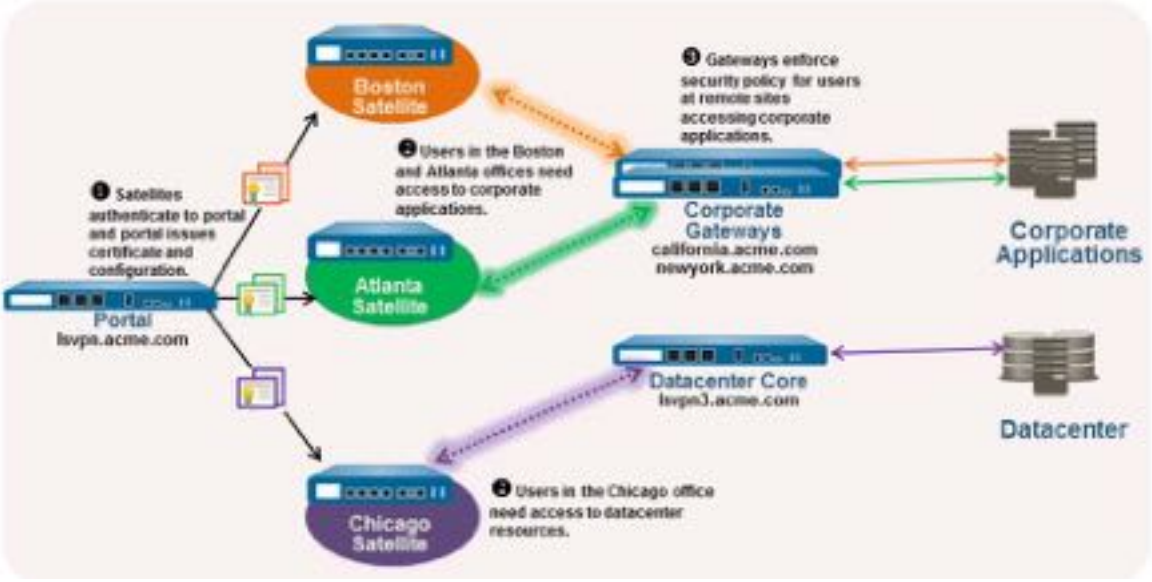


Use the following procedure to create one or more satellite configurations.

| Create a GlobalProtect Satellite Configuration | | |
|---|---|---|
| Step 1 | Specify the certificates required to enable satellites to participate in the LSVPN. | 1. Select **Network > GlobalProtect > Portals** and select the portal configuration for which you want to add a satellite configuration and then select the **Satellite Configuration** tab.<br>2. In the **Trusted Root CA** field, click **Add** and then select the CA certificate used to issue the gateway server certificates. The portal will deploy the root CA certificate you add here to all satellites as part of the configuration to enable the satellite to establish an SSL connection with the gateways. As a best practice, all of your gateways should use the same issuer.<br>If the root CA certificate used to issue your gateway server certificates is not on the portal, you can **Import** it now. See Enable SSL Between GlobalProtect LSVPN Components for details on how to import a root CA certificate.<br>3. Select the Root CA certificate that the portal will use to issue certificates to satellites upon successfully authenticating them from the **Issuing Certificate** drop-down. |
| Step 2 | Add a satellite configuration.<br>The satellite configuration specifies the GlobalProtect LSVPN configuration settings to deploy to the connecting satellites. You must define at least one satellite configuration. | In the Satellite Configuration section, click **Add** and enter a **Name** for the configuration.<br>If you plan to create multiple configurations, make sure the name you define for each is descriptive enough to allow you to distinguish them. |

| Create a GlobalProtect Satellite Configuration (Continued) | |
|---|---|
| Step 3 Specify which satellites to deploy this configuration to.<br><br>The portal uses the **Enrollment User/User Group** settings and/or **Devices** serial numbers to match a satellite to a configuration. Therefore, if you have multiple configurations, be sure to order them properly. As soon as the portal finds a match, it will deliver the configuration. Therefore, more specific configurations must precede more general ones. See Step 6 for instructions on ordering the list of satellite configurations. | Specify the match criteria for the satellite configuration as follows:<br>• To restrict this configuration to satellite devices with specific serial numbers, select the **Devices** tab, click **Add**, and enter serial number (you do not need to enter the satellite hostname; it will be automatically added when the satellite connects). Repeat this step for each satellite you want to receive this configuration.<br>• Select the **Enrollment User/User Group** tab, click **Add**, and then select the user or group you want to receive this configuration. Satellites that do not match on serial number will be required to authenticate as a user specified here (either an individual user or group member).<br><br>Before you can restrict the configuration to specific groups, you must Map Users to Groups. |
| Step 4 Specify the gateways that satellites with this configuration can establish VPN tunnels with.<br><br>Routes published by the gateway are installed on the satellite as static routes. The metric for the static route is 10x the routing priority. If you have more than one gateway, make sure to also set the routing priority to ensure that routes advertised by backup gateways have higher metrics compared to the same routes advertised by primary gateways. For example, if you set the routing priority for the primary gateway and backup gateway to 1 and 10 respectively, the satellite will use 10 as the metric for the primary gateway and 100 as the metric for the backup gateway. | 1. On the **Gateways** tab, click **Add**.<br>2. Enter a descriptive **Name** for the gateway. The name you enter here should match the name you defined when you configured the gateway and should be descriptive enough identify the location of the gateway.<br>3. Enter the FQDN or IP address of the interface where the gateway is configured in the **Gateways** field. The address you specify must exactly match the Common Name (CN) in the gateway server certificate.<br>4. (Optional) If you are adding two or more gateways to the configuration, the **Routing Priority** helps the satellite pick the preferred gateway. Enter a value in the range of 1-25, with lower numbers having the higher priority (that is, the gateway the satellite will connect to if all gateways are available). The satellite will multiply the routing priority by 10 to determine the routing metric. |
| Step 5 Save the satellite configuration. | 1. Click **OK** to save the satellite configuration.<br>2. If you want to add another satellite configuration, repeat Step 2 through Step 5. |
| Step 6 Arrange the satellite configurations so that the proper configuration is deployed to each satellite. | • To move a satellite configuration up on the list of configurations, select the configuration and click **Move Up**.<br>• To move a satellite configuration down on the list of configurations, select the configuration and click **Move Down**. |

 © Palo Alto Networks, Inc.

| Create a GlobalProtect Satellite Configuration (Continued) | |
|---|---|
| Step 7 Save the portal configuration. | 1. Click **OK** to save the settings and close the GlobalProtect Portal dialog.<br>2. **Commit** your changes. |

# Prepare the Satellite Device to Join the LSVPN

In order to participate in the LSVPN, the satellite devices require a minimal amount of configuration. Because the required configuration is minimal, you can pre-configure the devices before shipping them to your branch offices for installation.

| Prepare the Satellite Device to Join the GlobalProtect LSVPN | | |
|---|---|---|
| Step 1 | Configure a Layer 3 interface. | This is the physical interface the satellite will use to connect to the portal and the gateway. This interface must be in a zone that allows access outside of the local trust network. As a best practice, create a dedicated zone for VPN connections for visibility and control over traffic destined for the corporate gateways. |
| Step 2 | Configure the logical tunnel interface for the tunnel to use to establish VPN tunnels with the GlobalProtect gateways.<br> IP addresses are not required on the tunnel interface unless plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues. | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select an existing zone or create a separate zone for VPN tunnel traffic by clicking **New Zone** and defining a **Name** for new zone (for example *lsvpnsat*).<br>4. In the **Virtual Router** drop-down, select **default**.<br>5. (Optional) If you want to assign an IP address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 2.2.2.11/24.<br>6. To save the interface configuration, click **OK**. |

 © Palo Alto Networks, Inc.

| Prepare the Satellite Device to Join the GlobalProtect LSVPN (Continued) | | |
|---|---|---|
| Step 3 | If you generated the portal server certificate using a Root CA that is not trusted by the satellites (for example, if you used self-signed certificates), import the root CA certificate used to issue the portal server certificate.<br><br>The root CA certificate is required to enable the satellite device to establish the initial connection with the portal to obtain the LSVPN configuration. | 1. Download the CA certificate that was used to generate the portal server certificates. If you are using self-signed certificates, export the root CA certificate from the portal as follows:<br>a. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>b. Select the CA certificate, and click **Export**.<br>c. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (You do not need to export the private key.)<br>2. Import the root CA certificate you just exported onto each satellite as follows.<br>a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.<br>c. **Browse** to the **Certificate File** you downloaded from the CA.<br>d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.<br>e. Select the certificate you just imported on the **Device Certificates** tab to open it.<br>f. Select **Trusted Root CA** and then click **OK**. |
| Step 4 | Configure the IPSec tunnel configuration. | 1. Select **Network > IPSec Tunnels** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the IPSec configuration.<br>3. Select the **Tunnel Interface** you created for the satellite.<br>4. Select **GlobalProtect Satellite** as the **Type**.<br>5. Enter the IP address or FQDN of the portal as the **Portal Address**.<br>6. Select the Layer 3 **Interface** you configured for the satellite.<br>7. Select the **Local IP Address** to use on the selected interface. |

| Prepare the Satellite Device to Join the GlobalProtect LSVPN (Continued) | |
|---|---|
| Step 5 (Optional) Configure the satellite to publish local routes to the gateway.<br><br>Pushing routes to the gateway enables traffic to the subnets local to the satellite via the gateway. However, you must also configure the gateway to accept the routes as detailed in Step 8 in Configure the Gateway. | 1. To enable the satellite to push routes to the gateway, on the **Advanced** tab select **Publish all static and connected routes to Gateway**.<br><br>If you select this check box, the firewall will forward all static and connected routes from the satellite to the gateway. However, to prevent the creation of routing loops, the firewall will automatically apply route filters, such as the following:<br>• Default routes<br>• Routes within a virtual router other than the virtual router associated with the tunnel interface<br>• Routes using the tunnel interface<br>• Routes using the physical interface associated with the tunnel interface<br>2. (Optional) If you only want to push routes for specific subnets rather than all routes, click **Add** in the Subnet section and specify which subnet routes to publish. |
| Step 6 Save the satellite configuration. | 1. Click **OK** to save the IPSec tunnel settings.<br>2. Click **Commit**. |
| Step 7 If required, provide the credentials to allow the satellite to authenticate to the portal.<br><br>This step is only required if the portal was unable to find a serial number match in its configuration or if the serial number didn't work. In this case, the satellite will not be able to establish the tunnel with the gateway(s). | 1. Select **Network > IPSec Tunnels** and click the **Gateway Info** link in the Status column of the tunnel configuration you created for the LSVPN.<br>2. Click the **enter credentials** link in the **Portal Status** field and username and password required to authenticate the satellite to the portal.<br><br>After the portal successfully authenticates to the portal, it will receive its signed certificate and configuration, which it will use to connect to the gateway(s). You should see the tunnel establish and the **Status** change to **Active**.<br><br> |


© Palo Alto Networks, Inc.

# Verify the LSVPN Configuration

After configuring the portal, gateways, and satellite devices, verify that the satellites are able to connect to the portal and gateway and establish VPN tunnels with the gateway(s).

| Verify the LSVPN Configuration | | |
|---|---|---|
| Step 1 | Verify satellite connectivity with portal. | From the firewall hosting the portal, verify that satellites are successfully connecting by selecting **Network > GlobalProtect > Portal** and clicking **Satellite Info** in the Info column of the portal configuration entry.<br> |
| Step 1 | Verify satellite connectivity with the gateway(s). | On each firewall hosting a gateway, verify that satellites are able to establish VPN tunnels by selecting **Network > GlobalProtect > Gateways** and click **Satellite Info** in the Info column of the gateway configuration entry. Satellites that have successfully established tunnels with the gateway will display on the **Active Satellites** tab.<br> |
| Step 1 | Verify LSVPN tunnel status on the satellite. | On each firewall hosting a satellite, verify the tunnel status by selecting **Network > IPSec Tunnels** and verify active Status as indicated by a green icon.<br> |

# LSVPN Quick Configs

The following sections provide step-by-step instructions for configuring some common GlobalProtect LSVPN deployments:

- Basic LSVPN Configuration with Static Routing
- Advanced LSVPN Configuration with Dynamic Routing

 © Palo Alto Networks, Inc.

# Basic LSVPN Configuration with Static Routing

This quick config shows the fastest way to get up and running with LSVPN. In this example, a single firewall at the corporate headquarters site is configured as both a portal and a gateway. Satellite devices can be quickly and easily deployed with minimal configuration for optimized scalability.

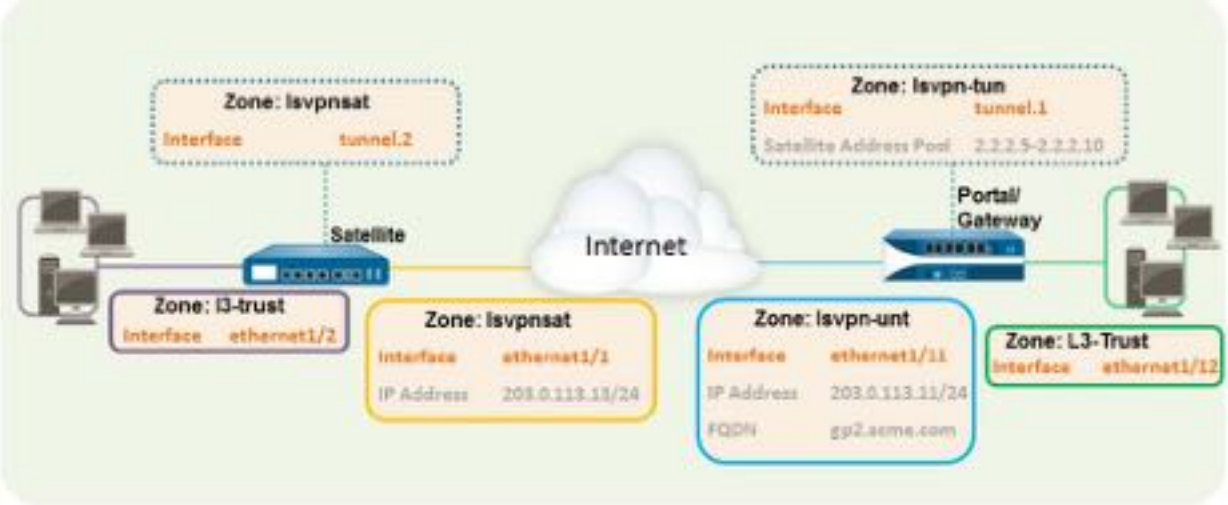


The following workflow shows the steps for setting up this basic configuration:

| Quick Config: Basic LSVPN with Static Routing | | |
|---|---|---|
| Step 1 | Configure a Layer 3 interface. | In this example, the Layer 3 interface on the portal/gateway requires the following configuration:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—lsvpn-unt<br>• **IPv4**—203.0.113.11/24 |
| Step 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.<br>To enable visibility into users and groups connecting over the VPN, enable User-ID in the zone where the VPN tunnels terminate. | In this example, the Tunnel interface on the portal/gateway requires the following configuration:<br>• **Interface**—tunnel.1<br>• **Security Zone**—lsvpn-tun |
| Step 3 | Create the security policy rule to enable traffic flow between the VPN zone where the tunnel terminates (lsvpn-tun) and the trust zone where the corporate applications reside (L3-Trust). | |

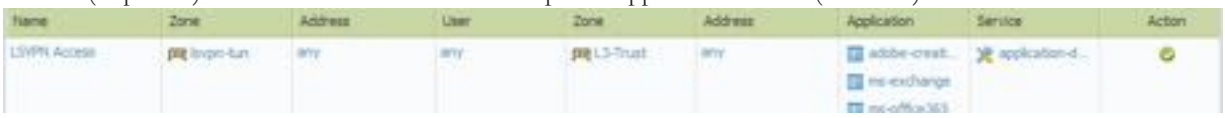

| Quick Config: Basic LSVPN with Static Routing (Continued) | | |
|---|---|---|
| Step 4 | Issue a self-signed server certificate for the portal/gateway.<br><br>The certificate subject name must match the FQDN or IP address of the Layer 3 interface you create for the portal/gateway. | 1. On the firewall hosting the portal, create the root CA certificate for issuing self-signed certificates for the GlobalProtect components. In this example, the root CA certificate, *lsvpn-CA*, will be used to issue the server certificate for the portal/gateway. In addition, the portal will use this root CA certificate to sign the CSRs from the satellite devices.<br>2. Generate server certificates for the GlobalProtect portal and gateway(s).<br><br>Because the portal and gateway will be on the same interface in this example, they can share a server certificate. In this example, the server certificate is named *lsvpnserver*. |
| Step 5 | Create a certificate profile. | In this example, the certificate profile, *lsvpn-profile*, references the root CA certificate *lsvpn-CA*. The gateway will use this certificate profile to authenticate satellites attempting to establish VPN tunnels. |
| Step 6 | Configure an authentication profile for the portal to use if the satellite serial number is not available. | 1. Create a server profile on the portal.<br>2. Create an authentication profile. In this example, the profile *lsvpn-sat* is used to authenticate satellites. |
| Step 7 | Configure the Gateway for LSVPN. | Select **Network > GlobalProtect > Gateways** and **Add** a configuration. This example requires the following gateway configuration:<br>• **Interface**—ethernet1/11<br>• **IP Address**—203.0.113.11/24<br>• **Server Certificate**—lsvpnserver<br>• **Certificate Profile**—lsvpn-profile<br>• **Tunnel Interface**—tunnel.1<br>• **Primary DNS**/**Secondary DNS**—4.2.2.1/4.2.2.2<br>• **IP Pool**—2.2.2.111-2.2.2.120<br>• **Access Route**—10.2.10.0/24 |
| Step 8 | Configure the Portal for LSVPN. | Select **Network > GlobalProtect > Portal** and **Add** a configuration. This example requires the following portal configuration:<br>• **Interface**—ethernet1/11<br>• **IP Address**—203.0.113.11/24<br>• **Server Certificate**—lsvpnserver<br>• **Authentication Profile**—lsvpn-sat |

 © Palo Alto Networks, Inc.

| Quick Config: Basic LSVPN with Static Routing (Continued) | |
|---|---|
| Step 9 Create a GlobalProtect Satellite Configuration. | On the **Satellite Configuration** tab in the portal configuration, **Add** a Satellite Configuration and a Trusted Root CA and specify the CA the portal will use to issue certificates for the satellites. In this example the required settings are as following:<br>• **Gateway**—203.0.113.11<br>• **Issuing Certificate**—lsvpn-CA<br>• **Trusted Root CA**—lsvpn-CA |
| Step 10 Prepare the Satellite Device to Join the LSVPN. | The satellite configuration in this example requires the following settings:<br>Interface Configuration<br>• Layer 3 interface—ethernet1/1, 203.0.113.13/24<br>• Tunnel interface—tunnel.2<br>• Zone—lsvpnsat<br>Root CA Certificate from Portal<br>• lsvpn-CA<br>IPSec Tunnel Configuration<br>• **Tunnel Interface**—tunnel.2<br>• **Portal Address**—203.0.113.11<br>• **Interface**—ethernet1/1<br>• **Local IP Address**—203.0.113.13/24<br>• **Publish all static and connected routes to Gateway**—enabled |

# Advanced LSVPN Configuration with Dynamic Routing

In larger LSVPN deployments with multiple gateways and many satellites, investing a little more time in the initial configuration to set up dynamic routing will simplify the maintenance of gateway configurations because access routes will update dynamically. The following example configuration shows how to extend the basic LSVPN configuration to configure OSPF as the dynamic routing protocol.

Setting up an LSVPN to use OSPF for dynamic routing requires the following additional steps on the gateways and the satellites:

- Manual assignment of IP addresses to tunnel interfaces on all gateways and satellites.
- Configuration of OSPF point-to-multipoint (P2MP) on the virtual router on all gateways and satellites. In addition, as part of the OSPF configuration on each gateway, you must manually define the tunnel IP address of each satellite as an OSPF neighbor. Similarly, on each satellite, you must manually define the tunnel IP address of each gateway as an OSPF neighbor.

Although dynamic routing requires additional setup during the initial configuration of the LSVPN, it reduces the maintenance tasks associated with keeping routes up to date as topology changes occur on your network.

The following figure shows an LSVPN dynamic routing configuration. This example shows how to configure OSPF as the dynamic routing protocol for the VPN.

For a basic setup of a LSVPN, follow the steps in Basic LSVPN Configuration with Static Routing. You can then complete the steps in the following workflow to extend the configuration to use dynamic routing rather than static routing.

 © Palo Alto Networks, Inc.

| Quick Config: LSVPN with Dynamic Routing | |
|---|---|
| Step 1 Add an IP address to the tunnel interface configuration on each gateway and each satellite.<br><br>**Important:**<br><br>The IP addresses you assign to the satellite tunnel interfaces must be on separate subnets from the IP addresses you assign to the gateway tunnel interfaces. In addition, the IP addresses you assign to satellites must not overlap with the designated IP pool defined in the gateway configuration or the devices will not be able to establish adjacencies. | Complete the following steps on each gateway and each satellite:<br>1. Select **Network > Interfaces > Tunnel** and select the tunnel configuration you created for the LSVPN to open the Tunnel Interface dialog.<br>If you have not yet created the tunnel interface, see Step 2 in Quick Config: Basic LSVPN with Static Routing.<br>2. On the **IPv4** tab, click **Add** and then enter an IP address and subnet mask. For example, to add an IP address for the gateway tunnel interface you would enter 2.2.2.100/24.<br>3. Click **OK** to save the configuration. |
| Step 2 Configure the dynamic routing protocol on the gateway.<br> 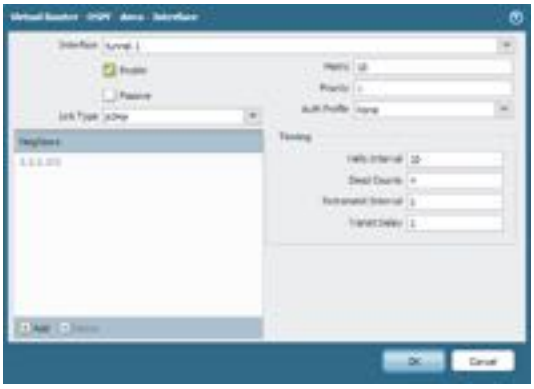 | To configure OSPF on the gateway:<br>1. Select **Network > Virtual Routers** and select the virtual router associated with your VPN interfaces.<br>2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.<br>3. If you are creating a new area, enter an **Area ID** on the **Type** tab.<br>4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.<br>5. Select **p2mp** as the **Link Type**.<br>6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each satellite device, for example 2.2.2.111.<br>7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.<br>8. Repeat this step each time you add a new satellite to the LSVPN. |
| Step 3 Configure the dynamic routing protocol on the satellite.<br> 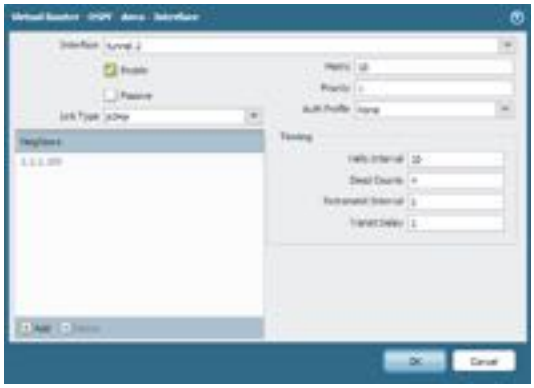 | To configure OSPF on the satellite:<br>1. Select **Network > Virtual Routers** and select the virtual router associated with your VPN interfaces.<br>2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.<br>3. If you are creating a new area, enter an **Area ID** on the **Type** tab.<br>4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.<br>5. Select **p2mp** as the **Link Type**.<br>6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each GlobalProtect gateway, for example 2.2.2.100.<br>7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.<br>8. Repeat this step each time you add a new gateway. |

**Quick Config: LSVPN with Dynamic Routing (Continued)**

Step 4 Verify that the gateways and satellites are able to form router adjacencies.

- On each satellite and each gateway, confirm that peer adjacencies have formed and that routing table entries have been created for the peers (that is, the satellites have routes to the gateways and the gateways have routes to the satellites). Select **Network > Virtual Router** and click the **More Runtime Stats** link for the virtual router you are using for the LSVPN. On the Routing tab, verify that the LSVPN peer has a route.

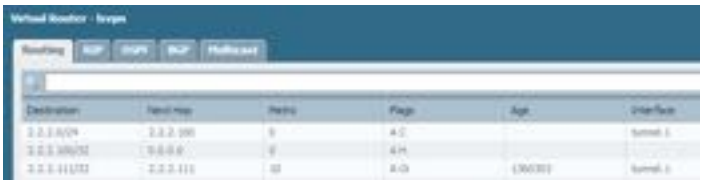

- On the **OSPF > Interface** tab, verify that the **Type** is **p2mp**.

- On the **OSPF > Neighbor** tab, verify that the firewalls hosting your gateways have established router adjacencies with the firewalls hosting your satellites and vice versa. Also verify that the **Status** is **Full**, indicating that full adjacencies have been established.

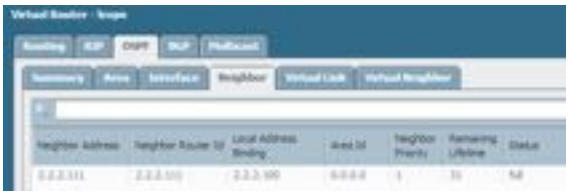

 © Palo Alto Networks, Inc.

# Networking

All Palo Alto Networks next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, enabling you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports. The following sections provide basic information on each type of deployment. For more detailed deployment information, refer to Designing Networks with Palo Alto Networks Firewalls and for information on route distribution, refer to Understanding Route Redistribution and Filtering.

The following topics describe how to integrate Palo Alto Networks next-generation firewalls into your network.

- Interface Deployments
- Configure a Virtual Router
- Configure Static Routes
- Configure RIP
- Configure OSPF
- Configure BGP

# Interface Deployments


- Virtual Wire Deployments
- Layer 2 Deployments
- Layer 3 Deployments
- Tap Mode Deployments


## Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together and should be used only when no switching or routing is needed.

A virtual wire deployment allows the following conveniences:

- Simplifies installation and configuration.
- Does not require any configuration changes to surrounding or adjacent network devices.

The "default-vwire" that is shipped as the factory default configuration, binds together Ethernet ports 1 and 2 and allows all untagged traffic. You can, however, use a virtual wire to connect any two ports and configure it to block or allow traffic based on the virtual LAN (VLAN) tags; the VLAN tag "0" indicates untagged traffic. You can also create multiple subinterfaces, add them into different zones and then classify traffic according to a VLAN tag, or a combination of a VLAN tag with IP classifiers (address, range, or subnet) to apply granular policy control for specific VLAN tags or for VLAN tags from a specific source IP address, range, or subnet.

**Figure: Virtual Wire Deployment**

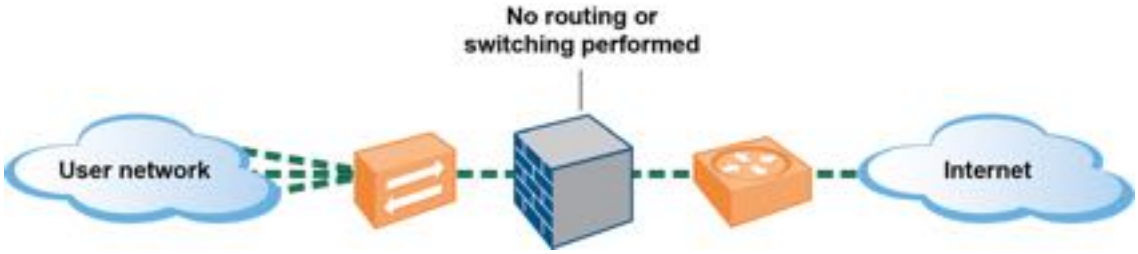


### Virtual Wire Subinterfaces

Virtual wire subinterfaces provide flexibility in enforcing distinct policies when you need to manage traffic from multiple customer networks. It allows you to separate and classify traffic into different zones (the zones can belong to separate virtual systems, if required) using the following criteria:

- **VLAN tags** —The example in Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only), shows an Internet Service Provider (ISP) using virtual wire subinterfaces with VLAN tags to separate traffic for two different customers.

 © Palo Alto Networks, Inc.

- **VLAN tags in conjunction with IP classifiers (address, range, or subnet)**— The following example shows an Internet Service Provider (ISP) with two separate virtual systems on a firewall that manages traffic from two different customers. On each virtual system, the example illustrates how virtual wire subinterfaces with VLAN tags and IP classifiers are used to classify traffic into separate zones and apply relevant policy for customers from each network.

| Virtual Wire Subinterface Workflow | |
|---|---|
| Step 1 | Configure two Ethernet interfaces as type virtual wire, and assign these interfaces to a virtual wire. |
| Step 2 | Create subinterfaces on the parent Virtual Wire to separate CustomerA and CustomerB traffic. Make sure that the VLAN tags defined on each pair of subinterfaces that are configured as virtual wire(s) are identical. This is essential because a virtual wire does not switch VLAN tags. |
| Step 3 | Create new subinterfaces and define IP classifiers. This task is optional and only required if you wish to add additional subinterfaces with IP classifiers for further managing traffic from a customer based on the combination of VLAN tags and a specific source IP address, range or subnet.<br><br>You can also use IP classifiers for managing untagged traffic. To do so, you must create a sub-interface with the vlan tag "0", and define sub-interface(s) with IP classifiers for managing untagged traffic using IP classifiers |

IP classification may only be used on the subinterfaces associated with one side of the virtual wire. The subinterfaces defined on the corresponding side of the virtual wire must use the same VLAN tag, but must not include an IP classifier.

**Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only)**

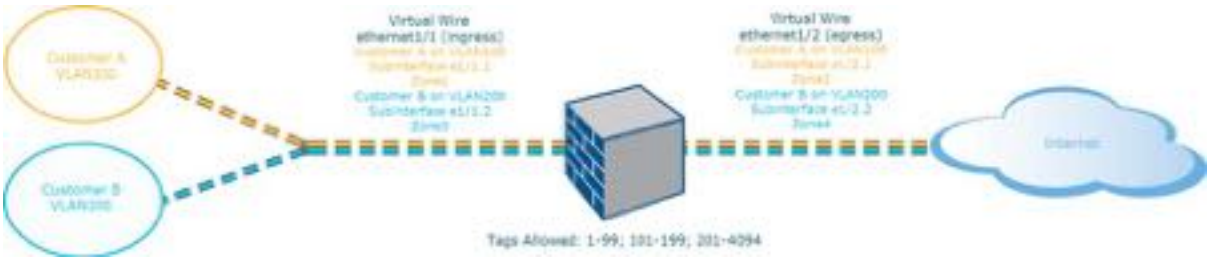


Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only) depicts CustomerA and CustomerB connected to the firewall through one physical interface, ethernet1/1, configured as a Virtual Wire; it is the ingress interface. A second physical interface, ethernet1/2, is also part of the Virtual Wire; it is the egress interface that provides access to the Internet. For CustomerA, you also have subinterfaces ethernet1/1.1 (ingress) and ethernet1/2.1 (egress). For CustomerB, you have the subinterface ethernet1/1.2 (ingress) and ethernet1/2.2 (egress). When configuring the subinterfaces, you must assign the appropriate VLAN tag and zone in order to apply policies for each customer. In this example, the policies for CustomerA are created between Zone1 and Zone2, and policies for CustomerB are created between Zone3 and Zone4.

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this example, a single subinterface matches the VLAN tag on the incoming packet, hence that subinterface is selected. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

The same VLAN tag must not be defined on the parent virtual wire interface and the subinterface. Verify that the VLAN tags defined on the Tag Allowed list of the parent virtual wire interface (**Network > Virtual Wires**) are not included on a subinterface.

Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers) depicts CustomerA and CustomerB connected to one physical firewall that has two virtual systems (vsys), in addition to the default virtual system (vsys1). Each virtual system is an independent virtual firewall that is managed separately for each customer. Each vsys has attached interfaces/subinterfaces and security zones that are managed independently.

**Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers)**

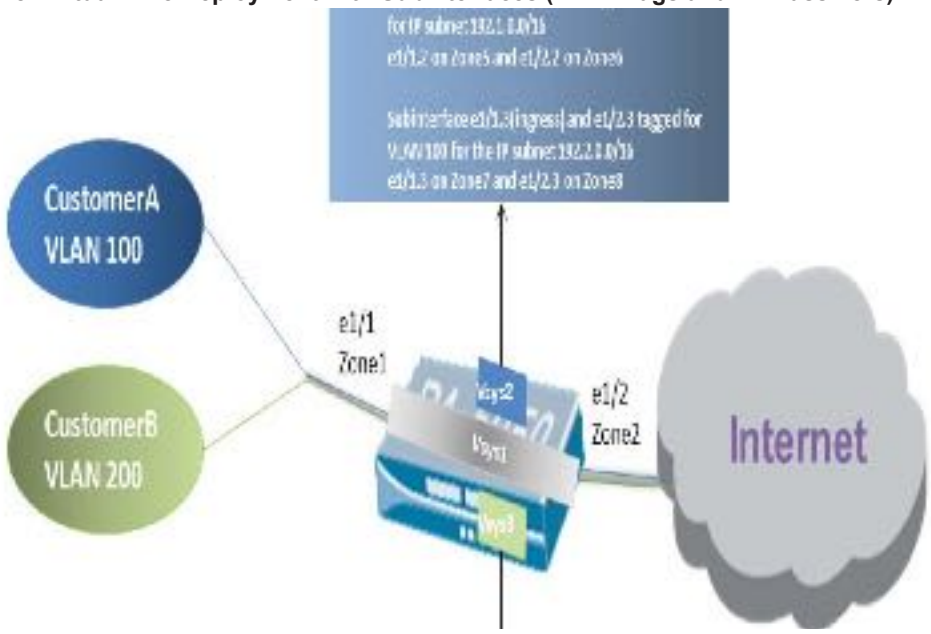


Vsys1 is set up to use the physical interfaces ethernet1/1 and ethernet1/2 as a virtual wire; ethernet1/1 is the ingress interface and ethernet1/2 is the egress interface that provides access to the Internet. This virtual wire is configured to accept all tagged and untagged traffic with the exception of VLAN tags 100 and 200 that are assigned to the subinterfaces.

CustomerA is managed on vsys2 and CustomerB is managed on vsys3. On vsys2 and vsys3, the following vwire subinterfaces are created with the appropriate VLAN tags and zones to enforce policy measures.

| Customer | Vsys | Vwire Subinterfaces | Zone | VLAN Tag | IP Classifier |
|---|---|---|---|---|---|
| A | 2 | e1/1.1 (ingress)<br>e1/2.1 (egress) | Zone3<br>Zone4 | 100<br>100 | None |
| | 2 | e1/1.2 (ingress)<br>e1/2.2 (egress) | Zone5<br>Zone6 | 100<br>100 | IP subnet<br>192.1.0.0/16 |
| | 2 | e1/1.3 (ingress)<br>e1/2.3 (egress) | Zone7<br>Zone8 | 100<br>100 | IP subnet<br>192.2.0.0/16 |
| B | 3 | e1/1.4 (ingress)<br>e1/2.4 (egress) | Zone9<br>Zone10 | 200<br>200 | None |

 © Palo Alto Networks, Inc.

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this case, for CustomerA, there are multiple subinterfaces that use the same VLAN tag. Hence, the firewall first narrows the classification to a subinterface based on the source IP address in the packet. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

For return-path traffic, the firewall compares the destination IP address as defined in the IP classifier on the customer-facing subinterface and selects the appropriate virtual wire to route traffic through the accurate subinterface.

The same VLAN tag must not be defined on the parent virtual wire interface and the subinterface. Verify that the VLAN tags defined on the Tag Allowed list of the parent virtual wire interface (**Network > Virtual Wires**) are not included on a subinterface.

## Layer 2 Deployments

In a Layer 2 deployment, the firewall provides switching between two or more networks. Each group of interfaces must be assigned to a VLAN object in order for the firewall to switch between them. The firewall will perform VLAN tag switching when layer 2 subinterfaces are attached to a common VLAN object. Choose this option when switching is required.

**Figure: Layer 2 Deployment**

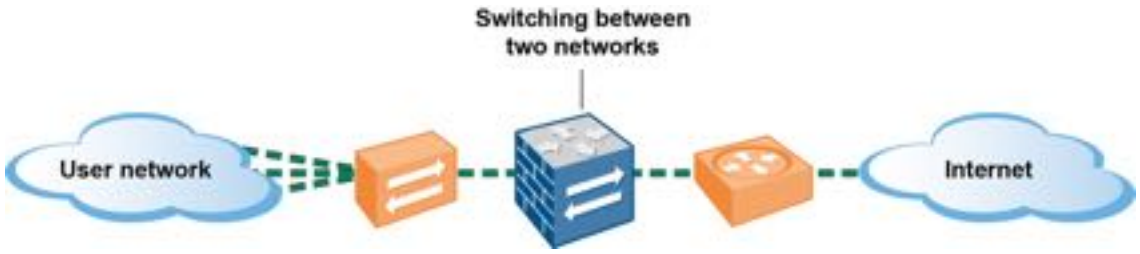


## Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between multiple ports. An IP address must be assigned to each interface and a virtual router must be defined to route the traffic. Choose this option when routing is required.

**Figure: Layer 3 Deployment**

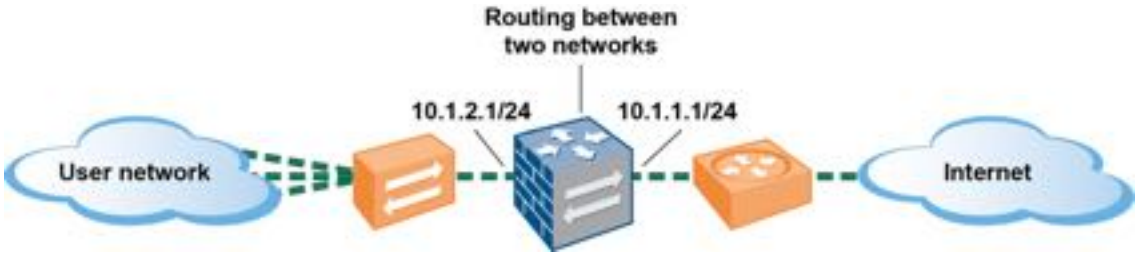

In addition, because the firewall must route traffic in a Layer 3 deployment, you must configure a virtual router. See Configure a Virtual Router.

### Point-to-Point Protocol over Ethernet Support

You can configure the firewall to be a Point-to-Point Protocol over Ethernet (PPPoE) termination point to support connectivity in a Digital Subscriber Line (DSL) environment where there is a DSL modem but no other PPPoE device to terminate the connection.

You can choose the PPPoE option and configure the associated settings when an interface is defined as a Layer 3 interface.

PPPoE is not supported in HA active/active mode.

### DHCP Client

You can configure the firewall interface to act as a DHCP client and receive a dynamically assigned IP address. The firewall also provides the capability to propagate settings received by the DHCP client interface into a DHCP server operating on the firewall. This is most commonly used to propagate DNS server settings from an Internet service provider to client machines operating on the network protected by the firewall.

DHCP client is not supported in HA active/active mode.

## Tap Mode Deployments

A network tap is a device that provides a way to access data flowing across a computer network. Tap mode deployment allows you to passively monitor traffic flows across a network by way of a switch SPAN or mirror port.

The SPAN or mirror port permits the copying of traffic from other ports on the switch. By dedicating an interface on the firewall as a tap mode interface and connecting it with a switch SPAN port, the switch SPAN port provides the firewall with the mirrored traffic. This provides application visibility within the network without being in the flow of network traffic.

When deployed in tap mode, the firewall is not able to take action, such as blocking traffic or applying QoS traffic control.

 © Palo Alto Networks, Inc.

# Configure a Virtual Router

The firewall uses virtual routers to obtain routes to other subnets by manually defining a route (static routes) or through participation in Layer 3 routing protocols (dynamic routes). The best routes obtained through these methods are used to populate the firewall's IP route table. When a packet is destined for a different subnet, the Virtual Router obtains the best route from this IP route table and forwards the packet to the next hop router defined in the table.

The Ethernet interfaces and VLAN interfaces defined on the firewall receive and forward the Layer 3 traffic. The destination zone is derived from the outgoing interface based on the forwarding criteria, and policy rules are consulted to identify the security policies to be applied. In addition to routing to other network devices, virtual routers can route to other virtual routers within the same firewall if a next hop is specified to point to another virtual router.

You can configure the virtual router to participate with dynamic routing protocols (BGP, OSPF, or RIP) as well as adding static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that are not shared between virtual routers, enabling you to configure different routing behaviors for different interfaces.

Each Layer 3 interface, loopback interface, and VLAN interface defined on the firewall must be associated with a virtual router. While each interface can belong to only one virtual router, multiple routing protocols and static routes can be configured for a virtual router. Regardless of the static routes and dynamic routing protocols configured for a virtual router, a common general configuration is required. The firewall uses Ethernet switching to reach other devices on the same IP subnet.

The following Layer 3 routing protocols are supported from Virtual Routers:

- RIP
- OSPF
- OSPFv3
- BGP

| **Define a Virtual Router General Configuration** | | |
|---|---|---|
| Step 1 | Gather the required information from your network administrator. | • Interfaces that you want to route<br>• Administrative distances for Static, OSFP internal, OSPF external, IBGP, EBGP and RIP |
| Step 2 | Create the virtual router and name it. | 1. Select **Network > Virtual Routers>.**<br>2. Click **Add** and enter a name for the virtual router.<br>3. Select interfaces to apply to the virtual router.<br>4. Click **OK**. |
| Step 3 | Select interfaces to apply to the virtual router. | 1. Click **Add** in the Interfaces box.<br>2. Select an already defined interface from the drop-down.<br>3. Repeat Step 2 for all interfaces that you want to add to the virtual router. |

© Palo Alto Networks, Inc.

| Define a Virtual Router General Configuration (Continued) | | |
|---|---|---|
| Step 4 | Set Administrative Distances for static and dynamic routing. | 1. Set Administrative Distances as required.<br>• **Static** — Range: 10-240, Default: 10<br>• **OSPF Internal** — Range: 10-240, Default: 30<br>• **OSPF External** — Range: 10-240, Default: 110<br>• **IBGP** —Range: 10-240, Default: 200<br>• **EBGP** —Range: 10-240, Default: 20<br>• **RIP** —Range: 10-240, Default: 120 |
| Step 5 | Save virtual router general settings. | Click **OK** to save your settings. |
| Step 6 | Commit your changes. | Click **Commit**. The device may take up to 90 seconds to save your changes. |


© Palo Alto Networks, Inc.

# Configure Static Routes

The following procedure shows how to integrate the firewall into the network using static routing.

| Set Up Interfaces and Zones | | |
|---|---|---|
| Step 1 | Configure a default route to your Internet router. | 1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog.<br>2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and enter the route in the **Destination** field (for example, 0.0.0.0/0).<br>3. Select the **IP Address** radio button in the **Next Hop** field and then enter the IP address and netmask for your Internet gateway (for example, 208.80.56.1).<br>4. Click **OK** twice to save the virtual router configuration. |
| Step 2 | Configure the external interface (the interface that connects to the Internet). | 1. Select **Network > Interfaces** and then select the interface you want to configure. In this example, we are configuring Ethernet1/3 as the external interface.<br>2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.<br>3. In the **Virtual Router** drop-down, select **default**.<br>4. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Untrust, and then click **OK**.<br>5. To assign an IP address to the interface, select the **IPv4** tab and **Static** radio button. Click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 208.80.56.100/24.<br>6. To enable you to ping the interface, select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. Enter a **Name** for the profile, select **Ping** and then click **OK**.<br>7. To save the interface configuration, click **OK**. |

© Palo Alto Networks, Inc.

| Set Up Interfaces and Zones (Continued) | | |
|---|---|---|
| Step 3 | Configure the interface that connects to your internal network.<br><br>In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT. See Configure NAT Policies for details. | 1. Select **Network > Interfaces** and select the interface you want to configure. In this example, we are configuring Ethernet1/4 as the internal interface.<br>2. Select **Layer3** from the **Interface Type** drop down.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Trust, and then click **OK**.<br>4. Select the same Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab and the **Static** radio button, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 4 | Configure the interface that connects to the DMZ. | 1. Select the interface you want to configure.<br>2. Select **Layer3** from the **Interface Type** drop down. In this example, we are configuring Ethernet1/13 as the DMZ interface.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example DMZ, and then click **OK**.<br>4. Select the Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab and the **Static** radio button, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 5 | Save the interface configuration. | Click **Commit**. |
| Step 6 | Cable the firewall. | Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment. |
| Step 7 | Verify that the interfaces are active.<br> 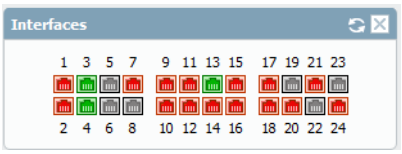 | From the web interface, select **Network > Interfaces** and verify that icon in the Link State column is green. You can also monitor link state from the **Interfaces** widget on the **Dashboard**. |

 © Palo Alto Networks, Inc.

# Configure RIP

RIP was designed for small IP networks and relies on hop count to determine routes; the best routes have the fewest number of hops. RIP is based on UDP and uses port 520 for route updates. By limiting routes to a maximum of 15 hops, the protocol helps prevent the development of routing loops, but also limits the supported network size. If more than 15 hops are required, traffic is not routed. RIP also can take longer to converge than OSPF and other routing protocols. The firewall supports RIP v2.

| **Configure RIP** | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Configure a Virtual Router for details. |
| Step 2 | Configure general RIP configuration settings. | 1. Select the **RIP** tab.<br>2. Select the **Enable** check box to enable the RIP protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through RIP. This is the recommended default setting.<br>4. De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through RIP. |
| Step 3 | Configure interfaces for the RIP protocol. | 1. Select the **Interfaces** subtab.<br>2. Select an interface from the drop-down in the Interface configuration box.<br>3. Select an already defined interface from the drop-down.<br>4. Select the **Enable** check box.<br>5. Select the **Advertise** check box to advertise a default route to RIP peers with the specified metric value.<br>6. You can optionally select a profile from the **Auth Profile** drop-down. See Step 5 for details.<br>7. Select normal, passive or send-only from the **Mode** drop-down.<br>8. Click **OK**. |
| Step 4 | Configure RIP timers. | 1. Select the **Timers** sub-tab.<br>2. Enter a value in the **Interval Seconds (sec)** box that defines the length of the timer interval in seconds. This duration is used for the remaining RIP timing fields. Default: 1. Range: 1 - 60.<br>3. Enter a value in the **Update Intervals** box that defines the number of intervals between route update announcements Default: 30. Range: 1 - 3600.<br>4. Enter a value in the **Delete Intervals** box that defines the number of intervals between the time that the route expires to its deletion Default: 180. Range: 1 - 3600.<br>5. Enter a value in the **Expire Intervals** box that defines the number of intervals between the time that the route was last updated to its expiration Default: 120. Range: 1 - 3600. |

© Palo Alto Networks, Inc.

| Configure RIP | |
|---|---|
| Step 5 (Optional) Configure Auth Profiles. | By default, the firewall does not use RIP authentication for the exchange between RIP neighbors. Optionally, you can configure RIP authentication between RIP neighbors by either a simple password or using MD5 authentication.<br>**Simple Password RIP authentication**<br>1. Select the **Auth Profiles** sub tab.<br>2. Click **Add**.<br>3. Enter a name for the authentication profile to authenticate RIP messages.<br>4. Select **Simple Password** as the **Password Type**.<br>5. Enter a simple password and then confirm.<br>**MD5 RIP authentication**<br>1. Select the **Auth Profiles** subtab.<br>2. Click **Add**.<br>3. Enter a name for the authentication profile to authenticate RIP messages.<br>4. Select **MD5** as the **Password Type**.<br>5. Click **Add**.<br>6. Enter one or more password entries, including:<br>• Key-ID Range 0-255<br>• Key<br>7. You can optionally select **Preferred** status.<br>8. Click **OK**. to specify the key to be used to authenticate outgoing message.<br>9. Click **OK** again in the **Virtual Router - RIP Auth Profile** configuration box. |

 © Palo Alto Networks, Inc.

# Configure OSPF

Open Shortest Path First (OSPF) is an interior gateway protocol (IGP) which is most often used to dynamically manage network routes in large enterprise network. It determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The information gathered from the LSAs is used to construct a topology map of the network. This topology map is shared across routers in the network and used to populate the IP routing table with available routes.

Changes in the network topology are detected dynamically and used to generate a new topology map within seconds. A shortest path tree is computed of each route. Metrics associated with each routing interface are used to calculate the best route. These can include distance, network throughput, link availability etc. Additionally, these metrics can be configured statically to direct the outcome of the OSPF topology map.

Palo Alto networks implementation of OSPF fully supports the following RFCs:

- RFC 2328 (for IPv4)
- RFC 5340 (for IPv6)

The following topics provide more information about the OSPF and procedures for configuring OSPF on the firewall:

▲ OSPF Concepts
▲ Configure OSPF
▲ Configure OSPFv3
▲ Configure OSPF Graceful Restart
▲ Confirm OSPF Operation

Also refer to How to Configure OSPF Tech Note.

## OSPF Concepts

The following topics introduce the OSPF concepts you will need to understand in order to configure the firewall to participate in an OSPF network:

▲ OSPFv3
▲ OSPF Neighbors
▲ OSPF Areas
▲ OSPF Router Types

## OSPFv3

OSPFv3 provides support for the OSPF routing protocol within an IPv6 network. As such, it provides support for IPv6 addresses and prefixes. It retains most of the structure and functions in OSPFv2 (for IPv4) with some minor changes. The following are some of the additions and changes to OSPFv3:

- **Support for multiple instances per link**—With OSPFv3 you can run multiple instances of the OSPF protocol over a single link. This is accomplished by assigning an OSPFv3 instance ID number. An interface that is assigned to an instance ID drops packets that contain a different ID.
- **Protocol Processing Per-link**—OSPFv3 operates per-link instead of per-IP-subnet as on OSPFv2.
- **Changes to Addressing**—IPv6 addresses are not present in OSPFv3 packets, except for LSA payloads within link state update packets. Neighboring routers are identified by the Router ID.
- **Authentication Changes**—OSPFv3 doesn't include any authentication capabilities. Configuring OSPFv3 on a firewall requires an authentication profile that specifies Encapsulating Security Payload (ESP) or IPv6 Authentication Header (AH).The re-keying procedure specified in RFC 4552 is not supported in this release.
- **Support for multiple instances per-link**—Each instance corresponds to an instance ID contained in the OSPFv3 packet header.
- **New LSA Types**—OSPFv3 supports two new LSA types: Link LSA and Intra Area Prefix LSA.

All additional changes are described in detail in RFC 5340.

## OSPF Neighbors

Two OSPF-enabled routers connected by a common network and in the same OSPF area that form a relationship are OSPF neighbors. The connection between these routers can be through a common broadcast domain or by a point-to-point connection. This connection is made through the exchange of hello OSPF protocol packets. These neighbor relationships are used to exchange routing updates between routers.

## OSPF Areas

OSPF operates within a single autonomous system (AS). Networks within this single AS however, can be divided into a number of Areas. By default, Area 0 is created. Area 0 can either function alone or act as the OSPF backbone for a larger number of Areas. Each OSPF area is named using a 32-bit identifier which in most cases written in the same dot-decimal notation as an IP4 address. For example, Area 0 is usually written as 0.0.0.0.

The topology of an area is maintained in its own link state database and is hidden from other areas which reduces the amount routing traffic required by OSPF. Topology is then shared in a summarized form between areas by a connecting router.

 © Palo Alto Networks, Inc.

### OSPF Area Types

**Backbone Area**—The backbone are (Area 0) is the core of an OSPF network. All other areas are connected to it and all traffic between areas must traverse it. All routing between areas is distributed through the backbone area. While all other OSPF areas must connect to the backbone area, this connection doesn't need to be direct and can be made through a virtual link.

**Normal OSPF Area**—In a normal OSPF area there are no restrictions; the area can carry all types of routes.

**Stub OSPF Area**—A stub area does not receive routes from other Autonomous Systems. Routing from the stub area is performed through the default route to the backbone area.

**NSSA Area**—The Not So Stubby Area (NSSA) is a type of stub area that can import external routes with some limited exceptions.

### OSPF Router Types

Within an OSPF area, routers are divided into the following categories.

**Internal Router**—A router with that only has OSPF neighbor relationships with devices in the same area.

**Area Border router (ABR)**—A router that has OSPF neighbor relationships with devices in multiple areas. ABRs gather topology information from their attached areas and distribute it to the backbone area.

**Backbone router**—A backbone router is any OSPF router that is attached to the OSPF backbone. Since ABRs are always connected to the backbone, they are always classified as backbone routers.

**Autonomous System Boundary Router (ASBR)**—An ASBR is a router that attaches to more than one routing protocol and exchanges routing information between them.

## Configure OSPF

OSPF determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The router keeps information about the links between it and the destination and can make highly efficient routing decisions. A cost is assigned to each router interface, and the best routes are determined to be those with the lowest costs, when summed over all the encountered outbound router interfaces and the interface receiving the LSA.

Hierarchical techniques are used to limit the number of routes that must be advertised and the associated LSAs. Because OSPF dynamically processes a considerable amount of route information, it has greater processor and memory requirements than does RIP.

| **Configure OSPF** | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Configure a Virtual Router for details. |

| Configure OSPF (Continued) | | |
|---|---|---|
| Step 2 | Configure general OSPF configuration settings. | 1. Select the **OSPF** tab.<br>2. Select the **Enable** check box to enable the OSPF protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through OSPF. This is the recommended default setting.<br>4. De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through OSPF. |


© Palo Alto Networks, Inc.

**Configure OSPF (Continued)**

| | |
|---|---|
| Step 3 Configure Areas Type for the OSFP protocol | 1. Select the **Areas** sub tab and click **Add**.<br>2. Enter an Area ID for the area in x.x.x.x format. This is the identifier that each neighbor must accept to be part of the same area.<br>3. Select the **Type** sub-tab.<br>4. Select one of the following from the area **Type** drop down box:<br>• **Normal** – There are no restrictions; the area can carry all types of routes.<br>• **Stub** – There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, configure the following:<br>– **Accept Summary** – Link state advertisements (LSA) are accepted from other areas. If this option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.<br>– **Advertise Default Route** – Default route LSAs will be included in advertisements to the stub area along with a configured metric value in the configured range: 1-255.<br>• **NSSA** (Not-So-Stubby Area) – The firewall can only leave the area by routes other than OSPF routes. If selected, configure **Accept Summary** and **Advertise Default Route** as described for **Stub**. If you select this option, configure the following:<br>– **Type** – Select either **Ext 1** or **Ext 2** route type to advertise the default LSA.<br>– **Ext Ranges** – Click **Add** in the section to enter ranges of external routes that you want to enable or suppress advertising for.<br>5. **Priority**—Enter the OSPF priority for this interface (0-255). It is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR) according to the OSPF protocol. When the value is zero, the router will not be elected as a DR or BDR.<br>• **Auth Profile**—Select a previously-defined authentication profile.<br>• **Timing**—It is recommended that you keep the default timing settings.<br>• **Neighbors**—For p2pmp interfaces, enter the neighbor IP address for all neighbors that are reachable through this interface.<br>6. Select **normal**, **passive** or **send-only** as the **Mode**.<br>7. Click **OK**. |

| Configure OSPF (Continued) | | |
|---|---|---|
| Step 4 | Configure Areas Range for the OSFP protocol | 1. Select the **Range** subtab.<br>2. Click **Add** to aggregate LSA destination addresses in the area into subnets.<br>3. **Advertise** or **Suppress** advertising LSAs that match the subnet, and click **OK**. Repeat to add additional ranges. |
| Step 5 | Configure Areas Interfaces for the OSFP protocol | 1. Select the **Interface** subtab.<br>2. Click **Add** and enter the following information for each interface to be included in the area, and click **OK**.<br>• **Interface**—Select an interface from the drop down box.<br>• **Enable**—Selecting this option causes the OSPF interface settings to take effect.<br>• **Passive**—Select the check box to if you do not want the OSPF interface to send or receive OSPF packets. Although OSPF packets are not sent or received if you choose this option, the interface is included in the LSA database.<br>• **Link type**—Choose **Broadcast** if you want all neighbors that are accessible through the interface to be discovered automatically by multicasting OSPF hello messages, such as an Ethernet interface. Choose **p2p** (point-to-point) to automatically discover the neighbor. Choose **p2mp** (point-to-multipoint) when neighbors must be defined manually. Defining neighbors manually is allowed only for p2mp mode.<br>• **Metric** — Enter an OSPF metric for this interface. Default: 10. Range: 0-65535.<br>• **Priority** — Enter an OSPF priority for this interface. This is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR). Default: 1. Range: 0 - 255. If zero is configured, the router will not be elected as a DR or BDR.<br>• **Auth Profile** — Select a previously-defined authentication profile.<br>• **Timing** — The following OSPF timing settings can be set here: Hello Interval, Dead Counts, Retransmit Interval and Transit Delay. Palo Alto Networks recommends that you retain the default timing settings.<br>• If **p2mp** is selected for **Link Type**, interfaces, enter the neighbor IP addresses for all neighbors that are reachable through this interface. |


© Palo Alto Networks, Inc.

| Configure OSPF (Continued) | |
|---|---|
| Step 6 Configure Areas Virtual Links. | 1. Select the **Virtual Link** sub tab.<br>2. Click **Add** and enter the following information for each virtual link to be included in the backbone area, and click **OK**:<br>• **Name** — Enter a name for the virtual link.<br>• **Neighbor ID** — Enter the router ID of the router (neighbor) on the other side of the virtual link.<br>• **Transit Area** — Enter the area ID of the transit area that physically contains the virtual link.<br>• **Enable** — Select to enable the virtual link.<br>• **Timing** — It is recommended that you keep the default timing settings.<br>• **Auth Profile** — Select a previously-defined authentication profile. |
| Step 7 (Optional) Configure Auth Profiles. | By default, the firewall does not use OSPF authentication for the exchange between OSPF neighbors. Optionally, you can configure OSPF authentication between OSPF neighbors by either a simple password or using MD5 authentication.<br>**Simple Password OSPF authentication**<br>1. Select the **Auth Profiles** subtab.<br>2. Click **Add**.<br>3. Enter a name for the authentication profile to authenticate OSPF messages.<br>4. Select **Simple Password** as the **Password Type**.<br>5. Enter a simple password and then confirm.<br>**MD5 OSPF authentication**<br>1. Select the **Auth Profiles** subtab.<br>2. Click **Add**.<br>3. Enter a name for the authentication profile to authenticate OSPF messages.<br>4. Select **MD5** as the **Password Type**.<br>5. Click **Add**.<br>6. Enter one or more password entries, including:<br>• Key-ID Range 0-255<br>• Key<br>• Select the **Preferred** option to specify that the key be used to authenticate outgoing messages.<br>7. Click **OK**.<br>8. Click **OK** again in the **Virtual Router - OSPF Auth Profile** configuration box. |

| Configure OSPF (Continued) | | |
|---|---|---|
| Step 8 | Configure Advanced OSPF options. | 1. Select the **Advanced** subtab.<br>2. Select the **RFC 1583 Compatibility** check box to assure compatibility with RFC 1583.<br>3. Configure a value for the **SPF Calculation Delay** (sec) timer.<br>This timer allows you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times.<br>4. Configure a value for the **LSA Interval (sec)** time, This timer specifies the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur. |

## Configure OSPFv3

| Configure OSPFv3 | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Configure a Virtual Router for details. |
| Step 2 | Configure general OSPF configuration settings. | 1. Select the **OSPF** tab.<br>2. Select the **Enable** check box to enable the OSPF protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through OSPF. This is the recommended default setting.<br>4. De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through OSPF. |
| Step 3 | Configure general OSPFv3 configuration settings. | 1. Select the **OSPFv3** tab.<br>2. Select the **Enable** check box to enable the OSPF protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through OSPFv3 This is the recommended default setting.<br>De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through OSPFv3. |


© Palo Alto Networks, Inc.

| Configure OSPFv3 (Continued) | | |
|---|---|---|
| Step 4 | Configure Auth Profile for the OSFPv3 protocol.<br><br>While OSPFv3 doesn't include any authentication capabilities of its own, instead, it relies entirely on IPsec to secure communications between neighbors. | When configuring an authentication profile you must use Encapsulating Security Payload (ESP) or IPv6 Authentication Header (AH).<br><br>**ESP OSPFv3 authentication**<br>1. Select the **Auth Profiles** subtab.<br>2. Click **Add**.<br>3. Enter a name for the authentication profile to authenticate OSPFv3 messages.<br>4. Specify a Security Policy Index (**SPI**). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a HEX value between 00000000 and FFFFFFFF.<br>5. Select **ESP** for **Protocol**.<br>6. Select a **Crypto Algorithm** from the drop down box.<br>You can enter none or one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5.<br>7. If a **Crypto Algorithm** other than none was selected, enter a value for **Key** and then confirm.<br><br>**AH OSPFv3 authentication**<br>1. Select the **Auth Profiles** subtab.<br>2. Click **Add**.<br>3. Enter a name for the authentication profile to authenticate OSPFv3 messages.<br>4. Specify a Security Policy Index (**SPI**). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a HEX value between 00000000 and FFFFFFFF.<br>5. Select **AH** for **Protocol**.<br>6. Select a **Crypto Algorithm** from the drop-down.<br>You must enter one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5.<br>7. Enter a value for **Key** and then confirm.<br>8. Click **OK**.<br>9. Click **OK** again in the Virtual Router - OSPF Auth Profile dialog. |

| Configure OSPFv3 (Continued) | |
|---|---|
| Step 5 Configure Areas Type for the OSFP protocol. | 1. Select the **Areas** subtab.<br>2. Click **Add**.<br>3. Enter an Area ID. This is the identifier that each neighbor must accept to be part of the same area.<br>4. Select the **General** sub-tab.<br>5. Select one of the following from the area **Type** drop-down:<br>• **Normal** – There are no restrictions; the area can carry all types of routes.<br>• **Stub** – There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, configure the following:<br>– **Accept Summary** – Link state advertisements (LSA) are accepted from other areas. If this option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.<br>– **Advertise Default Route** – Default route LSAs will be included in advertisements to the stub area along with a configured metric value in the configured range: 1-255.<br>• **NSSA** (Not-So-Stubby Area) – The firewall can only leave the area by routes other than OSPF routes. If selected, configure **Accept Summary** and **Advertise Default Route** as described for **Stub**. If you select this option, configure the following:<br>– **Type** – Select either **Ext 1** or **Ext 2** route type to advertise the default LSA.<br>– **Ext Ranges** – Click **Add** in the section to enter ranges of external routes that you want to enable or suppress advertising for. |
| Step 6 Associate an OSPFv3 authentication profile to an area or an interface. | **To an Area**<br>1. Select the **Areas** subtab.<br>2. Select an existing area from the table**.**<br>3. Select a previously defined **Authentication Profile** from the **Authentication** drop-down on the **General** subtab.<br>4. Click **OK**.<br>**To an Interface**<br>1. Select the **Areas** subtab.<br>2. Select an existing area from the table**.**<br>3. Select the **Interface** subtab and click **Add**.<br>4. Select the authentication profile you want to associate with the OSPF interface from the **Auth Profile** drop- down. |


© Palo Alto Networks, Inc.

| Configure OSPFv3 (Continued) | |
|---|---|
| Step 7 (Optional) Configure Export Rules | 1. Select the **Export** subtab.<br>2. Click **Add**.<br>3. Select the **Allow Redistribute Default Route** check box to permit redistribution of default routes through OSPFv3.<br>4. Select the name of a redistribution profile. The value must be an IP subnet or valid redistribution profile name.<br>5. Select a metric to apply for **New Path Type**.<br>6. Specify a **New Tag** for the matched route that has a 32-bit value.<br>7. Assign a metric for the new rule.<br>The value can be: 1 - 65535.<br>8. Click **OK**. |
| Step 8 Configure Advanced OSPFv3 options. | 1. Select the **Advanced** subtab.<br>2. Select the **Disable Transit Routing for SPF Calculation** check box if you want the firewall to participate in OSPF topology distribution without being used to forward transit traffic.<br>3. Configure a value for the **SPF Calculation Delay** (sec) timer.<br>This timer allows you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times.<br>4. Configure a value for the **LSA Interval (sec) time**, This timer specifies the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur.<br>5. Configure the Graceful Restart section as described in Configure OSPF Graceful Restart. |

## Configure OSPF Graceful Restart

OSPF Graceful Restart directs OSPF neighbors to continue using routes through a device during a short transition when it is out of service. This increases network stability by reducing the frequency of routing table reconfiguration and the related route flapping that can occur during short periodic down times.

For a Palo Alto Networks firewall this involves the following operations:

- **Firewall as a restarting device**—In a situation where the firewall will be down for a short period of time or is unavailable for short intervals, it sends Grace LSAs to its OSPF neighbors. The neighbors must be configured to run in Graceful Restart Helper mode. In Helper Mode, the neighbors receive the Grace LSAs that inform it that the firewall will perform a graceful restart within a specified period of time defined as the **Grace Period**. During the grace period, the neighbor continues to forward routes through the firewall and to send LSAs that announce routes through the firewall. If the firewall resumes operation before expiration

of the grace period, traffic forwarding will continue as before without network disruption. If the firewall does not resume operation after the grace period has expired, the neighbors will exit helper mode and resume normal operation which will involve reconfiguring the routing table to bypass the firewall.

- **Firewall as a Graceful Restart Helper—**In a situation where neighboring routers may be down for a short periods of time the firewall can be configured to operate in Graceful Restart Helper mode. If configured in this mode, the firewall will be configured with a **Max Neighbor Restart Time**. When the firewall receives the Grace LSAs from its OSFP neighbor, it will continue to route traffic to the neighbor and advertise routes through the neighbor until either the grace period or max neighbor restart time expires. If neither expires before the neighbor returns to service, traffic forwarding continues as before without network disruption. If either period expires before the neighbor returns to service, the firewall will exit helper mode and resume normal operation which will involve reconfiguring the routing table to bypass the neighbor.

| **Configure OSPF Graceful Restart** | |
|---|---|
| Step 1 | Select **Network > Virtual Routers** and select the virtual router you want to configure. |
| Step 2 | Select **OSPF > Advanced**. |
| Step 3 | Verify that the following check boxes are selected (they are enabled by default).<br>**Enable Graceful Restart, Enable Helper Mode, and Enable Strict LSA checking.**<br>All should remain selected unless required by your topology. |
| Step 4 | Configure a **Grace Period** in seconds. |
| Step 5 | Configure a **Max Neighbor Restart Time** in seconds. |

## Confirm OSPF Operation

Once an OSPF configuration has been committed, you can use any of the following operations to confirm that OSPF operating:

- View the Routing Table
- Confirm OSPF Adjacencies
- Confirm that OSPF Connections are Established

### View the Routing Table

By viewing the routing table, you can see whether OSPF routes have been established. The routing table is accessible from either the web interface or the CLI. If you are using the CLI, use the following commands:

- `show routing route`
- `show routing fib`

The following procedure describes how to use the web interface to view the route table.

 © Palo Alto Networks, Inc.

**View the Routing Table**

Step 1 Select **Network > Virtual Routers**

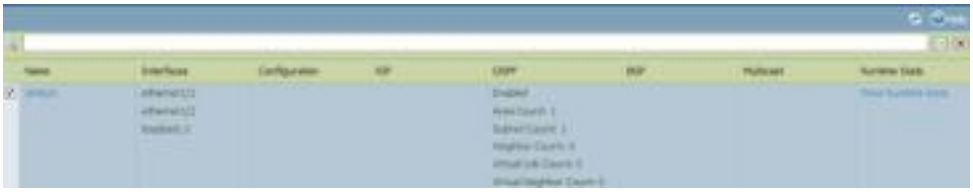

Step 2 Select the **Routing** tab and examine the **Flags** column of the routing table for routes that were learned by OSPF.

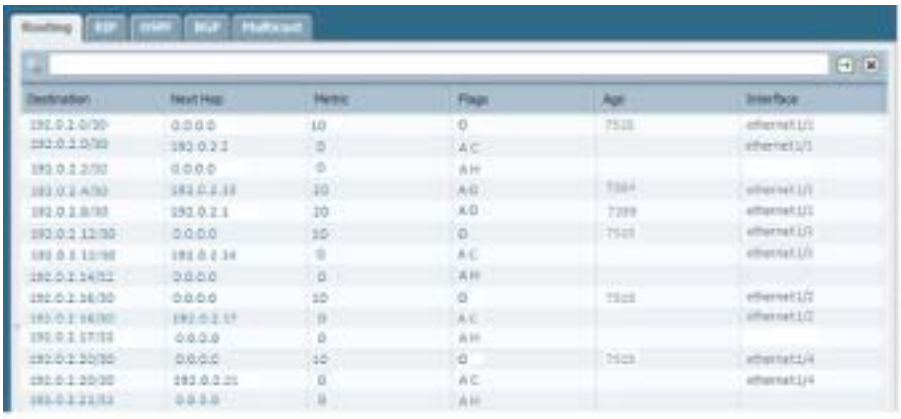

## Confirm OSPF Adjacencies

By viewing the **Neighbor** tab as described in the following procedure, you can confirm that OSPF adjacencies have been established.

**View the Neighbor Tab to Confirm OSPF Adjacencies**

Step 1 Select **Network > Virtual Routers**.

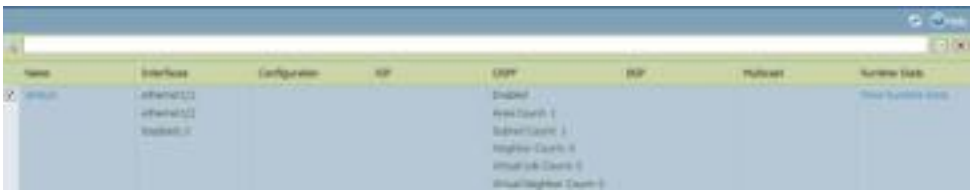

**View the Neighbor Tab to Confirm OSPF Adjacencies**

Step 2 Select **OSPF > Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established.

## Confirm that OSPF Connections are Established

By viewing the system log, you can confirm that OSPF connections have been established as described in the following procedure:

**Examine the System Log**

Step 1 Select **Monitor> System** and look for messages confirm that OSPF adjacencies have been established.

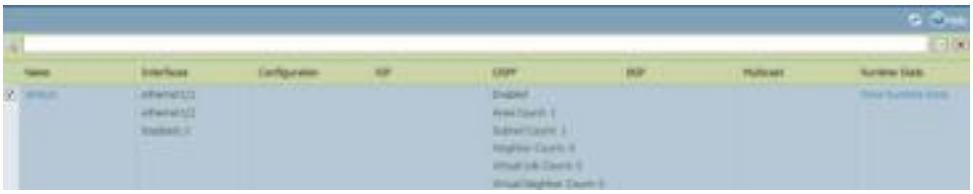

Step 2 Select **OSPF > Neighbor** sub tab and examine the **Status** column to determine if OSPF adjacencies have been established.

 © Palo Alto Networks, Inc.

# Configure BGP

The Border Gateway Protocol (BGP) is the primary Internet routing protocol. BGP determines network reachability based on IP prefixes that are available within autonomous systems (AS), where an AS is a set of IP prefixes that a network provider has designated to be part of a single routing policy.

In the routing process, connections are established between BGP peers (or neighbors). If a route is permitted by the policy, it is stored in the routing information base (RIB). Each time the local firewall RIB is updated, the firewall determines the optimal routes and sends an update to the external RIB, if export is enabled.

Conditional advertisement is used to control how BGP routes are advertised. The BGP routes must satisfy conditional advertisement rules before being advertised to peers.

BGP supports the specification of aggregates, which combine multiple routes into a single route. During the aggregation process, the first step is to find the corresponding aggregation rule by performing a longest match that compares the incoming route with the prefix values for other aggregation rules.

For more information on BGP, refer to How to Configure BGP Tech Note.

The firewall provides a complete BGP implementation that includes the following features:

- Specification of one BGP routing instance per virtual router.
- Routing policies based on route-map to control import, export and advertisement, prefix-based filtering, and address aggregation.
- Advanced BGP features that include route reflector, AS confederation, route flap dampening, and graceful restart.
- IGP-BGP interaction to inject routes to BGP using redistribution profiles.

BGP configuration consists of the following elements:

- Per-routing-instance settings, which include basic parameters such as local route ID and local AS and advanced options such as path selection, route reflector, AS confederation, route flap, and dampening profiles.
- Authentication profiles, which specify the MD5 authentication key for BGP connections.
- Peer group and neighbor settings, which include neighbor address and remote AS and advanced options such as neighbor attributes and connections.
- Routing policy, which specifies rule sets that peer groups and peers use to implement imports, exports, conditional advertisements, and address aggregation controls.

| **Configure BGP** | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Configure a Virtual Router for details. |

| Configure BGP (Continued) | | |
|---|---|---|
| Step 2 | Configure standard BGP configuration settings. | 1. Select the **BGP** tab.<br>2. Select the **Enable** check box to enable the BGP protocol.<br>3. Assign an IP address to the virtual router in the **Router ID** box.<br>4. Enter the number of the AS to which the virtual router belongs in the **AS Number** box, based on the router ID. Range: 1-4294967295 |
| Step 3 | Configure general BGP configuration settings. | 1. Select **BGP > General**.<br>2. Select the **Reject Default Route** check box to ignore any default routes that are advertised by BGP peers.<br>3. Select the **Install Route** check box to install BGP routes in the global routing table.<br>4. Select the **Aggregate MED** check box to enable route aggregation even when routes have different Multi-Exit Discriminator (MED) values.<br>5. Enter a value for the **Default Local Preference** that specifies a value than can be used to determine preferences among different paths.<br>6. Select one of the following values for the AS format for interoperability purposes:<br>• 2 Byte (default value)<br>• 4 Byte<br>7. Enable or disable each of the following values for **Path Selection**:<br>• **Always Compare MED** — Enable this comparison to choose paths from neighbors in different autonomous systems.<br>• **Deterministic MED Comparison** — Enable this comparison to choose between routes that are advertised by IBGP peers (BGP peers in the same autonomous system).<br>8. Click **Add** to include a new authentication profile and configure the following settings:<br>• **Profile Name**—Enter a name to identify the profile.<br>• **Secret/Confirm Secret**—Enter and confirm a passphrase for BGP peer communications. |


© Palo Alto Networks, Inc.

| Configure BGP (Continued) | |
|---|---|
| Step 4 Configure BGP Advanced settings (Optional) | 1. On the **Advanced** subtab, select **Graceful Restart** and configure the following timers:<br>• **Stale Route Time (sec)**—Specifies the length of time in seconds that a route can say in the stale state. Range: 1 - 3600 seconds. Default: 120 seconds.<br>• **Local Restart Time (sec)**—Specifies the length of time in seconds that the local device waits to restart. This value is advertised to peers. Range: 1 - 3600 seconds. Default: 120 seconds.<br>• **Max Peer Restart Time (sec)**—Specifies the maximum length of time in seconds that the local device accepts as a grave period restart time for peer devices. Range: 1 - 3600 seconds. Default: 120 seconds.<br>2. Specify an IPv4 identifier to represent the reflector cluster in the **Reflector Cluster ID** box.<br>3. Specify the identifier for the AS confederation to be presented as a single AS to external BGP peers in the **Confederation Member AS** box.<br>4. Click **Add** and enter the following information for each Dampening Profile that you want to configure, select **Enable**, and click **OK**:<br>• **Profile Name**—Enter a name to identify the profile.<br>• **Cutoff**—Specify a route withdrawal threshold above which a route advertisement is suppressed. Range: 0.0-1000.0. Default: 1.25.<br>• **Reuse**—Specify a route withdrawal threshold below which a suppressed route is used again. Range: 0.0-1000.0. Default: 5.<br>• **Max Hold Time (sec)**—Specify the maximum length of time in seconds that a route can be suppressed, regardless of how unstable it has been. Range: 0-3600 seconds. Default: 900 seconds.<br>• **Decay Half Life Reachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered reachable Range: 0-3600 seconds. Default: 300 seconds.<br>• **Decay Half Life Unreachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered unreachable. Range: 0 - 3600 seconds. Default: 300 seconds.<br>5. Click **OK**. |

| Configure BGP (Continued) | |
|---|---|
| Step 5 Configure the BGP peer group. | 1. Select the **Peer Group** subtab and click **Add**.<br>2. Enter a **Name** for the peer group and select **Enable**.<br>3. Select the **Aggregated Confed AS Path** check box to include a path to the configured aggregated confederation AS.<br>4. Select the **Soft Reset with Stored Info** check box to perform a soft reset of the firewall after updating the peer settings.<br>5. Specify the type of peer or group from the **Type** drop down box and configure the associated settings (see below in this table for descriptions of Import Next Hop and Export Next Hop).<br>• **IBGP—Export Next Hop:** Specify **Original** or **Use self**<br>• **EBGP Confed—Export Next Hop:** Specify **Original** or **Use self**<br>• **EBGP Confed—Export Next Hop:** Specify **Original** or **Use self**<br>• **EBGP—Import Next Hop:** Specify **Original** or **Use self**, **Export Next Hop:** Specify **Resolve** or **Use self**. Select **Remove Private AS** if you want to force BGP to remove private AS numbers.<br>6. Click **OK** to save. |
| Step 6 Configure Import and Export rules.<br>The import/export rules are used to import/export routes from/to other routers. For example, importing the default route from your Internet Service Provider. | 1. Select the **Import** tab and then click **Add** and enter a name in the **Rules** field and select the **Enable** check box.<br>2. Click **Add** and select the **Peer Group** to which the routes will be imported from.<br>3. Click the **Match** tab and define the options used to filter routing information. You can also define the Multi-Exit Discriminator (MED) value and a next hop value to routers or subnets for route filtering. The MED option is an external metric that lets neighbors know about the preferred path into an AS. A lower value is preferred over a higher value.<br>4. Click the **Action** tab and define the action that should occur (allow/deny) based on the filtering options defined in the **Match** tab. If **Deny** is selected, no further options need to be defined. If the **Allow** action is selected, define the other attributes.<br>5. Click the **Export** tab and define export attributes, which are similar to the **Import** settings, but are used to control route information that is exported from the firewall to neighbors.<br>6. Click **OK** to save. |


© Palo Alto Networks, Inc.

| Configure BGP (Continued) | | |
|---|---|---|
| Step 7 | Configure conditional advertising, which allows you to control what route to advertise in the event that a different route is not available in the local BGP routing table (LocRIB), indicating a peering or reachability failure. This is useful in cases where you want to try and force routes to one AS over another, for example if you have links to the Internet through multiple ISPs and you want traffic to be routed to one provider instead of the other unless there is a loss of connectivity to the preferred provider. | 1. Select the **Conditional Adv** tab, click **Add** and enter a name in the **Policy** field.<br>2. Select the **Enable** check box.<br>3. Click **Add** and in the **Used By** section enter the peer group(s) that will use the conditional advertisement policy.<br>4. Select the **Non Exist Filter** tab and define the network prefix(es) of the preferred route. This specifies the route that you want to advertise, if it is available in the local BGP routing table. If a prefix is going to be advertised and matches a Non Exist filter, the advertisement will be suppressed.<br>5. Select the **Advertise Filters** tab and define the prefix(es) of the route in the Local-RIB routing table that should be advertised in the event that the route in the non-exist filter is not available in the local routing table. If a prefix is going to be advertised and does not match a Non Exist filter, the advertisement will occur. |
| Step 8 | Configure aggregate options to summaries routes in the BGP configuration.<br><br>BGP route aggregation is used to control how BGP aggregates addresses. Each entry in the table results in one aggregate address being created. This will result in an aggregate entry in the routing table when at least one or more specific route matching the address specified is learned. | 1. Select the **Aggregate** tab, click **Add** and enter a name for the aggregate address.<br>2. In the **Prefix** field, enter the network prefix that will be the primary prefix for the aggregated prefixes.<br>3. Select the **Suppress Filters** tab and define the attributes that will cause the matched routes to be suppressed.<br>4. Select the **Advertise Filters** tab and define the attributes that will cause the matched routes to always be advertised to peers. |
| Step 9 | Configure redistribution rules.<br><br>This rule is used to redistribute host routes and unknown routes that are not on the local RIB to the peers routers. | 1. Select the Redist Rules tab and click Add.<br>2. In the **Name** field, enter an IP subnet or select a redistribution profile. You can also configure a new redistribution profile from the drop-down menu if needed.<br>3. Click the **Enable** check box to Enable the rule.<br>4. In the **Metric** field, enter the route metric that will be used for the rule.<br>5. In the **Set Origin** drop down, select incomplete, igp, or egp.<br>6. Optionally set MED, local preference, AS path limit and community values. |

 © Palo Alto Networks, Inc.

# PAN-OS® 6.0.13 Release Notes

**Revision Date: February 23, 2016**

Review important information about Palo Alto Networks PAN-OS 6.0 software, including new features introduced in this release, workarounds for open issues, and resolved issues. For the latest version of this release note, refer to the Palo Alto Networks Technical Documentation portal.


PAN-OS 6.0 Release Information . . . . . 3

Features Introduced in PAN-OS 6.0 . . . . . 4
- App-ID . . . . . 5
- Content Inspection Features . . . . . 6
- GlobalProtect Features . . . . . 10
- Management Features . . . . . 13
- Networking Features . . . . . 16
- Panorama Features . . . . . 18
- User-ID Features . . . . . 19
- Virtualization Features . . . . . 20

Changes to Default Behavior . . . . . 23

Associated Software Versions . . . . . 25

Known Issues . . . . . 26
- Known Issues in PAN-OS 6.0 . . . . . 26
- Known Issues Specific to VM-Series Firewalls . . . . . 28

PAN-OS 6.0.13 Addressed Issues . . . . . 31

PAN-OS 6.0.12 Addressed Issues . . . . . 33

PAN-OS 6.0.11 Addressed Issues . . . . . 35

PAN-OS 6.0.10 Addressed Issues . . . . . 39

PAN-OS 6.0.9 Addressed Issues . . . . . 41


© Palo Alto Networks, Inc.

PAN-OS 6.0.8 Addressed Issues . . . . . 45
PAN-OS 6.0.7 Addressed Issues . . . . . 49
PAN-OS 6.0.6 Addressed Issues . . . . . 53
PAN-OS 6.0.5-h3 Addressed Issues. . . . . 57
PAN-OS 6.0.5 Addressed Issues . . . . . 59
PAN-OS 6.0.4 Addressed Issues . . . . . 63
PAN-OS 6.0.3 Addressed Issues . . . . . 69
PAN-OS 6.0.2 Addressed Issues . . . . . 77
PAN-OS 6.0.1 Addressed Issues . . . . . 83
PAN-OS 6.0.0 Addressed Issues . . . . . 89
Getting Help . . . . . 97
Related Documentation . . . . . 97
Requesting Support . . . . . 98


© Palo Alto Networks, Inc.

# PAN-OS 6.0 Release Information

This release note provides important information about Palo Alto Networks® PAN-OS® 6.0 software, including an overview of new features introduced in this release and a list of known issues. For instructions on how to upgrade the firewall to PAN-OS 6.0 and configure the new features, refer to the New Features Guide.

For the most up-to-date information, refer to the online version of the PAN-OS 6.0 Release Notes on the Technical Documentation portal.


- Features Introduced in PAN-OS 6.0
- Changes to Default Behavior
- Associated Software Versions
- Known Issues

The client certificate used by PAN-OS and Panorama to authenticate to the PAN-DB cloud service, the WildFire cloud service, and the WF-500 appliance expired on January 21, 2016. The expiration results in an outage of these services.

To avoid an outage, either upgrade to content release version 550 (or a later version) or upgrade PAN-OS and Panorama instances running a PAN-OS or Panorama 6.0 release to PAN-OS (or Panorama) 6.0.12 or a later release.


- PAN-OS 6.0.13 Addressed Issues
- PAN-OS 6.0.12 Addressed Issues
- PAN-OS 6.0.11 Addressed Issues
- PAN-OS 6.0.10 Addressed Issues
- PAN-OS 6.0.9 Addressed Issues
- PAN-OS 6.0.8 Addressed Issues
- PAN-OS 6.0.7 Addressed Issues
- PAN-OS 6.0.6 Addressed Issues
- PAN-OS 6.0.5-h3 Addressed Issues
- PAN-OS 6.0.5 Addressed Issues
- PAN-OS 6.0.4 Addressed Issues
- PAN-OS 6.0.3 Addressed Issues
- PAN-OS 6.0.2 Addressed Issues
- PAN-OS 6.0.1 Addressed Issues
- PAN-OS 6.0.0 Addressed Issues
- Getting Help


© Palo Alto Networks, Inc.

# Features Introduced in PAN-OS 6.0

The following topics describe the new features introduced in the PAN-OS 6.0 release. This release requires content version 401 or later. For details on how to use the new features, refer to the New Features Guide.


- App-ID
- Content Inspection Features
- GlobalProtect Features
- Management Features
- Networking Features
- Panorama Features
- User-ID Features
- Virtualization Features



© Palo Alto Networks, Inc.

## App-ID

The following App-ID™ features are introduced in PAN-OS 6.0. For more details about these features and for instructions on configuring them, refer to Application Identification Features in the New Features Guide:

| New Application Identification Feature | Description |
|---|---|
| **Support for Hardware Security Modules** | You can now offload the certificate signing functions for SSL forward proxy, SSL inbound inspection, and the master key storage functions to a dedicated hardware security module (HSM) for enhanced key management security. This release supports the use of the following HSMs: SafeNet Luna SA and Thales nShield Connect. HSM support is generally required when FIPS 140-2 Level 3 protection for CA keys is required. The use of an HSM is supported on PA-7050, PA-3000 Series, PA-4000 Series, PA-5000 Series firewalls, VM-Series firewalls (VM 100, VM-200, VM-300, and VM-1000), and on the M-100 appliance. To configure a firewall to operate with an HSM, navigate to **Device > Setup > HSM**. |
| **Option to Disable SIP ALG** | By default, a Session Initiation Protocol (SIP) application-level gateway (ALG) performs NAT on the payload and opens dynamic pinholes for media ports. However, some SIP endpoints have NAT intelligence embedded in their clients. Because the firewall must not modify the signaling sessions in this case, you can now disable the SIP ALG functionality depending on the SIP applications in use in your environment. When SIP ALG is disabled, if App-ID determines that a session is SIP, the payload is not translated and open dynamic pinholes are not opened. To disable the SIP ALG, navigate to **Objects > Applications** and then customize the **ALG** option on the SIP application. |

© Palo Alto Networks, Inc.

## Content Inspection Features

The following Content Inspection features are introduced in PAN-OS 6.0. For more details about these features and for instructions on configuring them, refer to Content Inspection Features in the New Features Guide:

| New Content Inspection Feature | Description |
|---|---|
| **DNS Sinkholing** | This feature adds a new sinkhole action to the DNS signatures within the Anti-Spyware profile. Sinkholing enables the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. You can use this feature to identify infected hosts on the protected network using DNS traffic in situations where the firewall cannot see the infected client's DNS query (for example, when the firewall is north of the local DNS server). This feature can also be used to redirect malicious traffic to a honeypot or any other target host.<br>Enable sinkholing for the DNS signature collection by selecting the **sinkhole** action in the **DNS Signatures** tab of an Anti-Spyware profile and specifying an IPv4 and/or IPv6 address to use as the sinkhole (the default is the localhost, which will cut off communication). The sinkhole address can be an address of a live server or an unused address on your network. After enabling DNS sinkholing, you can identify infected clients by filtering the traffic logs or by building a custom report that checks for sessions to the IP address you defined as the sinkhole address. |
| **Extended Packet Capture** | A new extended-capture option has been added to Anti-Spyware and Vulnerability Protection profiles for rules and exceptions defined in the profile. Previously, when selecting packet capture, only the first trigger packet would be captured when a threat was detected in traffic matching the profiles. With the **extended-capture** option enabled, the firewall can capture from 1 to 50 packets, which provides much more context when analyzing the threat logs. To define the number of packets to capture, navigate to **Device > Setup > Content-ID** and then edit the **Threat Detection Settings** section. You can then view the extended packet captures from the threat logs (**Monitor > Logs > Threat**) by locating the log entry you want to investigate and then clicking the green arrow (Packet Capture) icon in the second column. |
| **Passive DNS** | This is an opt-in feature that enables the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis to improve threat intelligence and threat prevention capabilities. The data collected includes non-recursive DNS query and response packet payloads (such as payloads originating from the local recursive resolver, not individual clients). This information is used by the Palo Alto Networks threat research team to gain insight into malware propagation and evasion techniques that abuse the DNS system.<br>Information gathered through this data collection is used to improve accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control signatures, and the WildFire™ security service. Passive DNS monitoring is disabled by default but it is recommended that you enable it to facilitate enhanced threat intelligence. To enable this option select the **Enable Passive DNS Monitoring** check box on the **DNS Signatures** tab of the Anti-Spyware profile dialog for the Anti-Spyware profile that is attached to the security policy governing your DNS server's external DNS traffic. |

 © Palo Alto Networks, Inc.

| New Content Inspection Feature | Description |
|---|---|
| **URL Filtering Search Engine Cached Site Enhancement** | An enhancement has been made to the URL filtering engine such that URL filtering policies will also be applied when end-users attempt to view the Google and Internet Archive cached copies of websites. |
| **URL Filtering Translation Site Filtering Enhancement** | An enhancement has been made to the URL filtering engine such that URL filtering policies will also be applied to any URLs that are entered into translation sites such as Google Translate. This will ensure that website translation tools are not used to bypass URL filtering policies. |
| **URL Filtering Safe Search Enforcement** | This feature prevents users from viewing the search results unless the strictest safe search option is set in their browsers for and when searching using one of the top three search providers (Google, Bing, and Yahoo). If the strictest safe search option is not set in the browser, users will see a block page instructing them how to set the option for the given search provider. To enable this option, select the **Safe Search Enforcement** check box in the URL Filtering profile. Safe search will then be enforced whenever a user request matches a security policy rule with the corresponding URL Filtering profile attached.<br>Starting in PAN-OS 6.0.1, YouTube Safety Mode is supported as part of URL Filtering Safe Search Enforcement. YouTube Safety Mode ensures that YouTube search results only appear when the strictest safe search option is set in the user's browser.<br>Additionally, Safe Search Enforcement support for the Yandex search engine was added in content release version 446.<br>Finally, transparent safe search enforcement is supported with content release version 475 or later. |
| **WildFire Report Incorrect Verdict Option** | When viewing a WildFire analysis report for a sample that has been analyzed by WildFire, a new feature called **Report Incorrect Verdict** is now available in the report. This option enables you to resubmit the sample to the Palo Alto Networks threat team if you feel the verdict is a false positive or false negative. The threat team will perform further analysis on the sample to determine if it should be reclassified. If a file that was previously identified as malicious is determined to be benign (false positive), the signature for the file will be disabled in an upcoming antivirus signature update. Similarly if a file that was previously identified as benign is determined to be malicious (false negative), a new signature will be generated and distributed in the next content update. After the investigation is complete, you will receive an email (if provided on the submission form) notifying you of the outcome. |
| **WildFire Sandbox Operating Systems** | Microsoft Windows 7 32/bit has been added to the WildFire environment. When a dynamic analysis is performed on a file, it will be run in Windows XP and Windows 7. On a WF-500 WildFire appliance, you will need to select an image that will contain Windows XP or Windows 7 as well as a combination of other applications, such as different versions of Adobe Reader, and MS Office. |

© Palo Alto Networks, Inc.

| New Content Inspection Feature | Description |
|---|---|
| **Additional File Type Support Added to WildFire** | As part of the WildFire subscription, the following advanced file types are now supported: Microsoft Office (.doc/.docx, .xls/.xlsx, and .ppt/.pptx); Portable Document Format (.pdf); Java Applet (.jar and .class); and Android Application Package (.apk).<br>With a WildFire subscription, all listed file types can be submitted to WildFire from a PAN-OS firewall running PAN-OS 6.0 or later release versions using the WildFire web interface, a File Blocking profile in a security policy, or by manual upload to the WildFire portal at https://wildfire.paloaltonetworks.com. If you do not have a WildFire subscription, the firewall will forward only PE files. You can, however, manually upload any of the other supported file types to the WildFire portal.<br> The WF-500 WildFire appliance does not support APK file analysis. |
| **WildFire Analysis Report Enhancement** | The WildFire analysis report is now integrated with the logging features of the firewall and no longer requires a WildFire subscription. In addition, several new enhancements have been made to the report, including the ability to:<br>• Export the full report to a PDF by clicking **Download PDF**.<br>• Download the file sample that was analyzed.<br>• View the analysis results for each virtual environment in which the file was analyzed by clicking the corresponding **Virtual Machine** tab in the **Dynamic Analysis** section of the report. For example, click the **Virtual Machine 1** tab to view the Windows XP analysis results or click the **Virtual Machine 2** tab to see the analysis results for Windows 7. New virtual machine tabs will display as new environments are added. Each sandbox environment has its own configuration of applications and software used in the file analysis, such as which version of Adobe Reader, Flash, and MS Office are used in the file analysis on the specific virtual machine.<br>• View all processes or filter by an individual process.<br>• Re-submit the file sample to Palo Alto Networks for reevaluation if you think the file verdict (benign/malware) is incorrect by clicking **Report Incorrect Verdict**. |
| **WildFire Logging Update** | When a firewall is configured with a file blocking profile and security policy to forward files to WildFire for analysis, a WildFire subscription is no longer required to receive the WildFire Submissions logs on the firewall. A Subscription is still required to forward files to a WF-500 WildFire appliance and/or to forward the advanced file types that are now supported in PAN-OS 6.0 to either the WildFire cloud or a WildFire appliance. |
| **WildFire Report Integration** | Previously, in order to view the WildFire detailed report, you had to link to the report hosted on the WildFire cloud or the WildFire appliance. This report has now been integrated into the report on the firewall. Now, the WildFire Log Details report contains two tabs:<br>• **Details** tab-Shows the session details.<br>• **WildFire Analysis Report** tab-Shows the WildFire detailed report, which was previously hosted on the WildFire cloud or WildFire appliance.<br>Panorama™ no longer requires all managed firewalls to forward files to the same WildFire system so long as Panorama and the managed firewalls are running a PAN-OS 6.0 or later release version. |

 © Palo Alto Networks, Inc.

| New Content Inspection Feature | Description |
|---|---|
| **WildFire Submissions Log Forwarding** | Previously, you could only forward WildFire logs if you configured threat log forwarding for medium severity logs (which included WildFire logs with a malicious verdict) and/or informational severity logs (which included WildFire logs with a benign verdict). You can now configure the firewall to automatically forward WildFire Submissions logs independently of the threat log forwarding configuration. To enable WildFire Submissions log forwarding, configure the **WildFire Settings** in the log forwarding profile (**Objects > Log Forwarding**) to specify which WildFire logs (**Benign** and/or **Malicious** verdict) to forward to which log destinations (SNMP trap receiver, Panorama, syslog server, or email alert). |
| **WildFire API Enhancements** | The WildFire API syntax has changed for some web interface operations. For example, the file hash for sample files are now included in the logs, which you can use to query using the /get/report method. In addition, you can also now download WildFire PDF reports using the web interface and you can submit sample files to WildFire for the new files types that are now supported. For details on using the WildFire web interface in PAN-OS 6.0, refer to the WildFire Administrator's Guide. |

## GlobalProtect Features

The following GlobalProtect™ features are introduced in PAN-OS 6.0. For more details about these GlobalProtect features and for instructions on configuring them, refer to GlobalProtect Features in the New Features Guide.

| New GlobalProtect Feature | Description |
|---|---|
| **GlobalProtect Mobile Security Manager** | Provides, management, visibility, and automated configuration deployment for the mobile devices-either company provisioned or employee owned-on your network. With Mobile Security Manager you can create user- and/or HIP-based deployment policies that allow you to push application configurations (such as email and VPN configurations) to your employees' mobile devices. The Mobile Security Manager requires the devices it manages to check-in regularly to ensure that the device is in compliance, pushing new configuration to devices if the device status changes (for example, if the Mobile Security Manager determines it has an app with malware installed) or to push an updated policy to the device. You can also perform certain actions on a managed device from the Mobile Security Manager, such as locking the device, sounding an alarm to help locate the device, or even wiping a device that has been compromised. In addition, the GlobalProtect gateways can retrieve extended HIP report information for the devices managed by the Mobile Security Manager and use the information to enforce security policies for devices that connect to your network. The Mobile Security Manager runs on the GP-100 appliance. |
| **Agent Deployment Customization** | All GlobalProtect agent customization settings can now be set in the Windows registry (`HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\Settings\`) or the Mac global plist file (`/Library/Preferences/com.paloaltonetworks.GlobalProtect.settings.plist`). This enables deployment of GlobalProtect agent settings to client systems prior to their first connection to the GlobalProtect portal. For Windows clients, this also enables simplified deployment via group policy, as well as the option to automatically deploy the settings in the Windows Installer (MSIEXEC). Note that settings defined in the GlobalProtect portal client configuration take precedence over settings defined in the Windows Registry or the Mac plist.<br><br>In addition, the following new configuration options have been added to the **Network > GlobalProtect > Portals Agent** tab in the **Configs** dialog:<br>• **Show GlobalProtect icon—**Disabling this option makes GlobalProtect invisible to the end user by removing the icon from the system tray and preventing the user from having any interaction with the GlobalProtect agent user interface.<br>• **Allow user to change portal address—**Disabling this option prevents users from manually changing the portal address pushed to the client in the portal configuration file.<br>• **Allow user to continue if portal server certificate is invalid—**Disabling this option prevents users from continuing if there is a warning screen indicating a man-in-the-middle (MITM) attack, |


© Palo Alto Networks, Inc.

| New GlobalProtect Feature | Description |
|---|---|
| **Transparent One-Time Password (OTP) Support** | To simplify the GlobalProtect user authentication process and make it more transparent for the end user when authenticating to the portal and the gateway, the portal now includes settings for modifying the default authentication behavior on a per-client configuration basis. The following **Authentication Modifier** settings are now available on the **General** tab when configuring a client configuration on the portal (**Network > GlobalProtect > Portals**):<br>• **Cookie authentication for config refresh—**Enables the agent to use an encrypted cookie to authenticate to the portal when refreshing a configuration that has already been cached (the user will always be required to authenticate to the portal for the initial configuration download and upon cookie expiration). This simplifies the authentication process for end users because they will no longer be required to log in to both the portal and the gateway in succession or enter multiple OTPs for authenticating to each. If this option is enabled, the portal will instead use the cookie to authenticate to the user.<br> The cookie will never be used for authentication to the gateway.<br>• **Different password for external gateway—**Disables the forwarding of credentials to some or all gateways, enabling the gateway to immediately prompt for its own set of credentials. This option speeds up the authentication process when the portal and the gateway require different credentials (either different OTPs or different login credentials entirely). Or, you can choose to use a different password on manual gateways only. With this option, the portal will forward credentials to automatic gateways but not to manual gateways, allowing you to have the same security on your portals and automatic gateways, while requiring a second factor OTP or a different password for access to those gateways that provide access to your most sensitive resources. |
| **HIP Profile Support for Client DLP Products** | The GlobalProtect agent by default now collects vendor-specific data about whether data loss prevention (DLP) software is installed and/or enabled on Windows hosts. DLP software is used to prevent sensitive corporate information from leaving the corporate network or from being stored on a potentially insecure device. Because this information is now collected from Windows host systems, you can include DLP as matching criteria for the host information profiles (HIPs) you create, thereby enabling you to use DLP compliance as criteria for your security policies. |
| **Agent Update Control** | The client configurations delivered by the GlobalProtect portal now have two additional options for controlling when users can upgrade the GlobalProtect agent: **disable** and **manual**. This means there are now four options for the **Agent Upgrade** field on the **Agent** tab in the Configs dialog of the portal configuration:<br>• By default, the portal will **prompt** the agent to upgrade whenever a new version is available.<br>• To enable automatic upgrades that do not require user interaction, select **transparent**.<br>• (New option) To prevent agent upgrades, select **disable**.<br>• (New option) To allow end users to initiate agent upgrades, select **manual**. In this case, the user would select the **Check Version** option in the agent to determine if there is a new agent version and then upgrade if desired. |

| New GlobalProtect Feature | Description |
|---|---|
| **Certificate Authentication Enforcement** | Enhancements have been made to how client certificate authentication is enforced in various scenarios as follows:<br>• If a certificate profile is configured on the GlobalProtect portal, the client must present a certificate in order to connect. This means that certificates must be deployed to the end clients before their initial portal connection.<br>• If the certificate profile specifies a **Username Field**, the certificate presented by the client must contain a username in order to connect. Furthermore, if both an authentication profile and a certificate profile with the **Username Field** are configured, the end user will be forced to use the username from the certificate to log in for authentication against the configured authentication profile.<br>• For agents configured with the pre-logon connect method, if the new **Cookie authentication for config refresh** setting is enabled, you no longer need to configure a certificate profile for pre-logon authentication; in this case the portal will use the cookie to authenticate the client prior to user logon. Note, however, that you must still configure a certificate profile on the gateway to enable establishment of the VPN tunnel. |


© Palo Alto Networks, Inc.

## Management Features

The following Management features are introduced in PAN-OS 6.0. For more details about these features and for instructions on configuring them, refer to Management Features in the New Features Guide.

| New Management Feature | Description |
|---|---|
| **New Documentation Portal** | The all new web-based Palo Alto Networks Technical Documentation portal is now available at https://www.paloaltonetworks.com/documentation.html. The new content you will find here features concepts and workflows that focus on enabling start-to-finish network security solutions. The documentation suite now covers the full Palo Alto Networks product and feature set, including the following new guides introduced in this release:<br>• PAN-OS 6.0 New Features Guide<br>• The new and improved PAN-OS Administrator's Guide<br>• GlobalProtect Administrator's Guide<br>• VM-Series Deployment Guide<br>As of PAN-OS 6.0.4, this site has been enhanced so that all of the new workflow-based content is available in HTML with PDF downloads available on demand. This new format simplifies access to the documentation by allowing you to search for answers to specific questions, browse by product, or narrow the results by search facet. For details on how to get the most out of the new portal, watch the video. |
| **Commit Improvement** | The commit operation in Panorama and in PAN-OS has been enhanced to allow configuration edits during a commit. For example, if two administrators are logged in to the same firewall and the first administrator performs a commit, the second administrator can make updates to the configuration during the commit. This enhancement does not, however, allow multiple administrators to commit simultaneously. |
| **Content Delivery Network (CDN)/Update Server Verification** | Palo Alto Networks will begin transitioning to use of a CDN for delivery of software and content updates to devices running Panorama/PAN-OS 6.0 or later release version. This new process will provide a secure and distributed infrastructure to improve the software update process and license installation/verification for customers around the world.<br>In addition, a new option has been added to strengthen all secure communication between firewalls/Panorama and the Palo Alto Networks update servers. To enable this option, select **Device** (or **Panorama**) **> Setup > Services** and select the **Verify Update Server Identity** check box in the **Services** dialog. When enabled, Panorama or PAN-OS will not perform a software/content download unless the update server has an SSL certificate signed by a trusted authority. This will help ensure that software updates and license verification will only be performed from Palo Alto Networks and will eliminate the possibility of man-in-the-middle attacks. |

© Palo Alto Networks, Inc.

| New Management Feature | Description |
|---|---|
| **Enumeration of Rules within a Rulebase** | The **Policies** tab on the web interface includes a new column for displaying rule numbers. Each rule is automatically numbered and the ordering adjusts as rules are moved or reordered. When filtering rules to find rules that match the specified filter(s), each rule is listed with its number in the context of the complete set of rules in the rulebase and its place in the evaluation order.<br><br>In Panorama, pre-rules and post-rules are independently numbered. When rules are pushed from Panorama to a managed firewall, the rule numbering incorporates hierarchy in pre-rules, device rules, and post-rules within a rulebase and reflects the rule sequence and its evaluation order. |
| **Enhancements in Reports** | The reporting enhancements in this release include support for creating group activity reports and the ability to disable predefined reports.<br><br>To generate reports for user groups on the firewall, select **Monitor > PDF Reports > User Activity Report**, select Type: **Group** and pick the group for which to generate the report. Because Panorama does not have the user to group mapping information, you cannot generate group activity reports on Panorama.<br><br>To disable one or more predefined reports, select **Device** (or **Panorama**) **> Setup > Management**. Edit the Logging and Reporting Settings section and select the **Log Export and Reporting** tab and clear the check boxes for each report you want to disable. |
| **CLI Find Command** | The new `find command` CLI command helps you find a command when you don't know where to start looking in the hierarchy. This new command—which is available in all CLI modes—has two forms:<br>• You can use `find command` alone to display the entire command hierarchy in the current command mode.<br>Or<br>• You can use `find command` with the `keyword` argument to locate all commands that have the specified keyword. For example, to find all configure mode commands with the `username` keyword, you would enter the following:<br>`admin@mgmt-ui-4060# find command keyword username`<br>`set deviceconfig system log-export-schedule <name> protocol ftp username <value>`<br>`set deviceconfig system log-export-schedule <name> protocol scp username <value>`<br>`set deviceconfig setting wildfire session-info-select exclude-username <yes\|no>`<br>`set mgt-config password-complexity block-username-inclusion <yes\|no>`<br>`set network interface ethernet <name> layer3 pppoe username <value>`<br>`set shared certificate-profile <name> username-field` |


© Palo Alto Networks, Inc.

| New Management Feature | Description |
|---|---|
| **Support for Syslog over TCP and SSL** | The firewall and Panorama now support using TCP or SSL (default is UDP) for reliable and secure transport of logs to an external syslog server. SSLv3 and TLSv1 are supported and the default SSL port is 6514. To separate individual syslog messages in a TCP stream, the delimiter formats available are LF- Line Feed (BSD format, the default), and Message Length (IETF format). If the syslog server requires client authentication, you can configure the firewall/Panorama to use a certificate for secure communication. The option to mark a **Certificate for Secure Syslog** as available in **Device** (or **Panorama**) **> Certificate Management > Certificates > Device Certificates**.<br><br>To monitor and alert you to a connection failure, a system log of high severity is generated when the firewall or Panorama is disconnected from the syslog server; an SNMP trap is also generated. |
| **Support for Color-Coded Tags** | Tags allow you to group objects using keywords/phrases and color (optional) to visually distinguish objects. You can apply tags to address objects, address groups (static and dynamic), zones, services, service groups, and policy rules. Use the **Objects > Tags** tab to create a tag, assign tag color, or to delete, rename and clone tags. When tagged, the keyword can be used to sort or filter objects. |
| **SNMP Resource Monitoring Extensions** | All Palo Alto Networks firewalls support SNMP objects that provide resource utilization and failure reporting, including alerts for hardware failure, the insertion and removal of the power supply, disks, or system fans, and monitoring of resource utilization for high session use or load on the gateway.<br><br>The HOST-RESOURCES-MIB includes objects for monitoring the following:<br>• Memory and disk usage on the management plane. Swap utilization tracks how the swap space is used to store all the running programs that aren't being actively used to free up memory.<br>• Packet buffers on the dataplane.<br><br>The Enterprise MIBs include objects for monitoring the following:<br>• GlobalProtect gateway utilization in the PAN-Common MIB to monitor the total number of current active tunnels on the firewall device and the percentage on current tunnel utilization that is based on the number of active tunnels and the maximum number of tunnels allowed.<br>• Session utilization in the PAN-Common MIB for each virtual system to track the active sessions, and the percentage of current session utilization based on session limit configured for the virtual system. |
| **Enhancement in the Syslog Header** | You can now choose the format of the hostname field in the syslog header. The syslog header can display one of the following: FQDN (hostname and domain name), hostname, the IPv4 address, or the IPv6 address of the sending device. Configure this option in the **Send Hostname in Syslog** drop-down on the **Device** (or **Panorama**) **> Setup > Management > Logging and Reporting Settings** section. The drop-down provides the following choices: **FQDN**, **Hostname**, **IPv4 address**, **IPv6 address**, or **None**; select **None** to leave the hostname field in the syslog header empty. |

## Networking Features

The following Networking features are introduced in PAN-OS 6.0. For more details about these features and for instructions on configuring them, refer to Networking Features in the New Features Guide.

| New Networking Feature | Description |
|---|---|
| **OSPFv3 Support** | OSPFv3 provides support for the OSPF routing protocol within an IPv6 network. OSPFv3 offers similar structure and functionality to OSPFv2 (for IPv4). OSPPv3 and OSPFv2 can be run concurrently on the same firewall in a dual stack configuration. To configure a firewall to operate with OSPFv3, navigate to **Network > Virtual Routers** and then add or edit a virtual router configuration. |
| **OSPF Graceful Restart** | Reduces high availability (HA) active/passive failover times by allowing OSPF neighbors to continue using routes through a device during a short transition while out of service. This increases network stability by reducing the network outage experienced when OSPF adjacencies are reestablished. To configure a firewall to operate with OSPF Graceful Restart, navigate to **Network > Virtual Routers** and edit any OSPF or OSPFv3 virtual router configuration. |
| **IKE PKI Certificate Authentication for IPSec Site-to-Site VPNs** | With this release authentication security has been enhanced over previous releases that only supported pre-shared-key authentication. To configure a firewall for IKE PKI certificate authentication, navigate to **Network > IKE Gateways** and then edit any IKE gateway configuration to select **Certificate** as the **Authentication** type and then specify which certificates to use for **Local Identification** and **Peer Identification**. |
| **Decryption Port Mirror** | Provides the ability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as an RSA NetWitness or Solera Security Analytics Platform appliance—for archiving and analysis. This feature is necessary for organizations that require comprehensive data capture for forensic and historical purposes or data leak prevention (DLP) functionality. To enable use of this feature, you must download and install a free license. You can then configure a decryption profile (**Objects > Decryption Profile**) to enable forwarding to the decrypt mirror interface. Decryption port mirroring is available on the PA-7050, PA-5000 Series and PA-3000 Series platforms only. |
| **TLS 1.2 Decryption** | Previous releases of PAN-OS only supported TLS version 1.0. This release provides the Palo Alto Networks firewall with the ability to decrypt inbound sessions and forward proxy sessions that negotiate with TLS 1.2. |
| **Increase Jumbo Frame Size** | The maximum transmission unit (MTU) size has been increased to provide compatibility with equipment from other vendors. Jumbo frames must specifically be enabled at the global level. Once enabled, the default MTU size for all Layer 3 interfaces (the **Global MTU**) is set to a value of 9192 bytes, but can be configured for any value in the range of 512 - 9216 bytes. Because this value is applied globally, any other value required must be explicitly configured on a per-interface basis. To configure jumbo frame support, navigate to **Device > Setup > Session** and then select the **Enable Jumbo Frame** check box in the Session Settings dialog. After enabling jumbo frame support, you can customize the default MTU to be used globally and the minimum NAT64 MTU. |
| **Remove TCP Timestamp** | A new option has been added to the Zone Protection profile to enable you to strip the TCP timestamp from the TCP header. |


© Palo Alto Networks, Inc.

| New Networking Feature | Description |
|---|---|
| **IPv6 Neighbor Discovery Table Capacity Increase** | In previous versions of the PAN-OS software, the Neighbor Discovery (ND) table was smaller than the ARP table for IPv4. With this release the IPv6 ND table size has been increased to account for larger IPv6 networks and implementation of v4/v6, dual-stack configurations. |
| **Enhanced Use for Address Objects** | A Layer 3 interface on the firewall can now use an address object in lieu of an IP address. Using an address object allows you to separate the object from its value/IP address. The address object (**Objects > Addresses**) only has to be defined once and then it can be referenced multiple places in policy configuration. When the IP address(es) that was defined for the address object changes, you can edit the address object and the change in value is automatically inherited by all instances where the address object is used.<br><br>For improved scalability, Panorama templates also allow you to add or select an address object when configuring a Layer 3 interface on the managed firewalls. Previously, when configuring a Layer 3 interface in a template you could not define a unique IP address for each managed firewall. With this new feature a template can reference an address object; the value for the address object can either be defined locally on the firewall or it can be defined as a shared object or as a device group object on Panorama. Defining the address object on the firewall allows you to configure a unique IP address for each managed device. To prevent a commit failure, you must create the address object on each firewall before pushing the template to the managed firewalls.<br><br>Panorama templates support address objects in the following locations:<br>• IP addresses for Layer 3 interfaces<br>• Virtual address of the virtual IP address in a high availability (HA) active/active configuration<br>• Service routes<br>• NAT policy for source translation<br>• GlobalProtect portal and gateway IP addresses<br>• GlobalProtect satellite and site-to-site IP addresses<br>• IKE Gateway local address<br>• Multicast static and candidate IP addresses<br>• BGP peer group local addresses<br>• Hardware Security Module client IP addresses |
| **Consolidation of Timers Used in a High Availability (HA) Setup** | High availability (HA) timers are used to detect a firewall failure and trigger a failover. To reduce the complexity in configuring HA timers, three profiles have been added: **Recommended**, **Aggressive** and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.<br><br>Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to accommodate your network requirements.<br><br>On upgrade, the current/existing HA settings are saved to the **Advanced** profile. If you prefer to load the preset values for the **Recommended** or **Aggressive** profiles, use the **HA Timer Settings** option in the **Election Settings** section of the **Device > High Availability** tab. |

## Panorama Features

The following Panorama™ features are introduced in PAN-OS 6.0. For more details about these features and for instructions on configuring them, refer to Panorama Features in the New Features Guide:

| New Virtualization Feature | Description |
|---|---|
| **Log Forwarding from Panorama** | Panorama now allows for forwarding of aggregated logs, email notifications, and SNMP traps to external servers. Forwarding logs from Panorama reduces the load on the firewalls and provides a reliable and streamlined approach to combine and forward logs/SNMP traps/email notifications to external destinations. To facilitate this change the Panorama interface has been changed.<br>See the Panorama Log Forwarding to an External Destination Per Platform table in the New Features Guide to configure log forwarding from Panorama: |
| **Scheduling Dynamic Updates from Panorama** | Dynamic updates for Application and Threats, WildFire, Antivirus, and URL Database can be scheduled using the **Schedules** link on the **Panorama > Device Deployment > Dynamic Updates** tab.<br>The frequency of the updates, and the option to only download or to download and install updates to all managed devices and managed collectors using Panorama is configurable. |
| **Support for PAN-DB and BrightCloud Databases** | In deployments where both PAN-DB and BrightCloud databases are used concurrently for URL filtering, Panorama provides the capability to create shared policies and push the policies to devices running different databases. When a mismatch occurs between the URL database vendor configured on Panorama and what is configured on the device, the device now maps and auto-migrates URL categories and URL profiles so that the policies are relevant for the database enabled on the device. |
| **Enhanced Use for Address Objects** | See Enhanced Use for Address Objects feature in the Networking Features section. |


© Palo Alto Networks, Inc.

## User-ID Features

The following User-ID™ features are introduced in PAN-OS 6.0. For more details about these features and for instructions on configuring them, refer to User-ID Features in the New Features Guide:

| New User Identification Feature | Description |
|---|---|
| **Windows Server 2012/2012 R2 Support**<br><br>(Supported in PAN-OS 6.0.3 and later release versions) | You can now install the Windows-based User-ID agent on a system running the Windows Server 2012/2012 R2 operating system and you can now install the Terminal Server agent on a system running the Microsoft Windows Server 2012 R2 operating system. The User-ID agent that runs on the firewall can also integrate with Windows Server 2012/2012 R2 for user mapping and group mapping. The procedures to install the User-ID agent and Terminal Server agent and to configure user mapping are the same as for servers supported in earlier releases (Windows Server 2003 and 2008). |
| **User-ID Integration with Syslog** | In environments with existing network services that authenticate users, such as wireless controllers, 802.1x devices, Apple Open Directory servers, or other network access control (NAC) mechanisms, the firewall can now listen for syslog messages from those services so that the User-ID agent (either the Windows agent or the agentless user mapping feature on the firewall) can extract the authentication events from the log. Syslog filters that you define allow User-ID to parse the messages and extract the IP addresses and usernames of users who successfully authenticated to the external service and add the information to the IP address to username mappings it maintains. Previously you had to use the XML API in order to integrate with these external devices and extract this information. |
| **Increased User-ID Active Users Limit** | To support organizations with a large number of active users requiring IP address to username mapping for policy enforcement, the active user limit has been increased on the high-end firewall platforms based on the memory capacity of the individual platforms.<br>The following list summarizes User-ID active limits on all Palo Alto Networks next-generation firewall platforms:<br>PA-7050 and PA-5060: 256,000<br>PA-5050 and PA-5020: 28,000<br>PA-4000 Series, PA-3000 Series, PA-2000 Series, PA-500, and PA-200: 64,000 |
| **API Enhancement to Support Custom Terminal Service Solutions** | The Terminal Services agent enables IP address to username mapping for users on Windows-based terminal servers by assigning a unique port range to individual users. The User-ID XML API has been extended to enable you to create scripts that allow for the same type of IP address to username mappings on non-Windows based terminal services. Specifically, the existing XML API calls that notify the firewall of IP address to username mappings have been extended to facilitate multi-user host log in and log off by including port mappings in addition to IP addresses. |

© Palo Alto Networks, Inc.

## Virtualization Features

The following virtualization features are introduced in PAN-OS 6.0. For more details about these features and for instructions on configuring them, refer to Virtualization Features in the New Features Guide:

| New Virtualization Feature | Description |
|---|---|
| **Virtual Machine (VM) Monitoring Agent** | In previous releases, you had to use external scripts and the XML API on the firewall to furnish information about virtual machine IP addresses to the firewall. This feature allows you to dynamically retrieve changes in your virtualized environment without making configuration changes on the firewall.<br>The Windows User-ID agent and the firewall can now be configured to proactively monitor the supported VM Sources: VMware ESXi server (4.1 and 5.0) and the VMware vCenter Server. The firewall supports up to 10 sources and the Windows User-ID agent supports up to 100 sources. In order to enable VM monitoring, the VM guest on the supported VM Sources must have VMware Tools installed and running. When you configure VM Information Sources on the **Device > VM Information Sources** tab, the firewall can access and retrieve the IP addresses of virtual machines (newly provisioned or modified) on the monitored sources. For each VM, you can also tag a predefined list of attributes. Each of these attributes can be used as match criterion in the new dynamic address groups feature and then referenced in policies. See Dynamic Address Groups for more details. |


© Palo Alto Networks, Inc.

| New Virtualization Feature | Description |
|---|---|
| **Dynamic Address Groups** | The Dynamic Address Groups feature allows you to dynamically update the network address of an object referenced in policy. Unlike a static address group where you specify the network address of a host, the members of a dynamic address group are populated using the match criteria that you define. The match criteria uses logical `and` or `or` operators; each host that you want to add to the dynamic address group must contain the tag that is defined in the match criteria.<br><br>Tags can be defined directly on the firewall or on Panorama or they can be dynamically defined using the XML API and registered with the firewall. When an IP address and the corresponding tag (one or more) are registered, each dynamic group evaluates the tags and updates the list of members in its group. This capability allows you to add/ remove/update the members of a group without making any changes on the firewall; a commit is not required for policy to take effect when you provision new hosts or decommission existing hosts on your network. The dynamic address group will use the tags to populate its members. All matching tags-both static and dynamic-are evaluated to populate the group. The difference between static and dynamic tags is that statically added tags are part of the configuration on the firewall, and dynamically added tags are part of the runtime configuration.<br><br>Because the members of the group are included based on the tags you match against, if an object has tags that match the specified criteria for the dynamic address group, the object will become a group member and the corresponding policy will be applied. For example, when provisioning a new server in a datacenter, you could tag the server's IP address as HTTP and DMZ. Using Dynamic Address Groups and the security rules on the firewall, the newly provisioned server would automatically be included in the existing policies to:<br>• Allow incoming HTTP connections to all machines tagged as HTTP.<br>• Allow all authenticated users to access all DMZ tagged machines through SSH and the web management console for the server.<br><br>If you have a virtual environment with VMware, instead of using scripts to call the XML API, you can use the **Virtual Machine (VM) Monitoring Agent** to retrieve information (network address and corresponding tags) on new servers/guests deployed on the monitored hosts/servers. To learn about this feature, see Virtual Machine (VM) Monitoring Agent. |
| **VM-Series on Citrix SDX** | The VM-Series firewall is now supported on the Citrix SDX hardware platform running Citrix XenServer version 6.0.2 or later. Deploying the VM-Series firewall (one or more instances) on the SDX server provides the ability to consolidate the NetScaler VPX appliance and the Palo Alto Networks VM-Series firewall on the same physical platform and protect north-south and/or east-west traffic on your network. For a list of supported SDX platforms, deployment examples, and installation instructions, refer to Set Up a VM-Series Firewall on the Citrix SDX Server in the VM-Series Deployment Guide. |

| New Virtualization Feature | Description |
|---|---|
| **Support for the VM-Series NSX Edition Firewall** | Traditionally, the lack of context between the security and virtual environments made implementing security policies a challenge at the data center. To meet the security challenges in the software defined datacenter, Palo Alto Networks and VMware introduce a joint solution to safely enable all datacenter traffic, including intra-server virtual machine communications.<br><br>NSX, the VMware Networking and Security platform that is designed for the software-defined datacenter, automates the process of deploying and provisioning the VM-Series firewall as a service (also called Security Virtual Machine) on ESXi servers.<br><br>Refer to the table of components in the Support for the VM-Series NSX Edition Firewall section of the New Features Guide for information about the components of this joint Palo Alto Networks and VMware solution.<br><br>To automate the process of deploying the VM-Series firewall, Panorama registers the VM-Series firewall as a service on the NSX Manager. The NSX Manager then deploys one VM-Series firewall on every ESXi host in a cluster. When a new ESXi host is added to the cluster, the NSX Manager automatically deploys a new firewall. Because each instance of the NSX edition of the VM-Series firewall is directly hooked in to the hypervisor, the firewall is seamlessly placed in the data path and can see all traffic that flows through the ESXi host.<br><br>The NSX Manager and the VM-Series firewall work in concert to enforce security; each provides a set of traffic management rules that are applied to the traffic on each ESXi host. The first set of rules determines which traffic to direct to the VM-Series firewall on each host, and the second set of rules determines how the firewall must process-allow, deny, inspect, and constrain- the application for enabling it safely on your network. The tight integration of the NSX Manager, Panorama and the firewalls in this solution augments the ability to enforce consistent security policies using dynamic address groups. The NSX Manager updates Panorama of all changes in the virtual environment, and Panorama pushes the updates to the firewalls. All policy rules that reference these dynamic address groups are updated to reflect the changes in the virtualized environment and security policies are consistently applied to secure all network resources. To learn more about this solution refer to information about how to Set Up a VM-Series NSX Edition Firewall in the VM-Series Deployment Guide. |


© Palo Alto Networks, Inc.

# Changes to Default Behavior

The following is a list of changes to default behavior in PAN-OS 6.0:

- Added a new SSL root certificate in WF-500 6.0.12 for secure communication with PAN-OS and Panorama.

  December 1, 2015 is the expiration date for the previous certificate. You should update WF-500 appliances running release 6.0.11 or earlier releases to WF-500 6.0.12 or a later release to ensure successful communication with PAN-OS and Panorama after that date.

- After upgrading from a PAN-OS 5.0 release to a PAN-OS 6.0 release, interface errors are reported by SNMP monitoring tools. Although interface errors existed in PAN-OS 5.0 and earlier releases, firewalls running PAN-OS 5.0 releases did not display those interface errors using SNMP. Instead, SNMP monitoring tools displayed a value of `zero` for interface errors. A fix was added in PAN-OS 6.0 so that an SNMP monitoring tool displays interface errors with the correct values. Refer to the following KB article for more detail on this change: https://live.paloaltonetworks.com/docs/DOC-6864.

- If the Safe Search Enforcement option is enabled and the firewall has content update **422** or later installed, Safe Search enforcement will also apply to YouTube searches. With the Safe Search for YouTube, users will be able to search YouTube for videos, but if YouTube Safety Mode is not enabled, the user will not be able to watch any videos. The Safety Mode icon is located at the bottom of the YouTube.com page and will either show **Safety On** or **Safety Off**.

- The output for the `show session` command is modified to include tracker session information. A new line shows `tracker stage firewall` with the following possible values: `TCP RST-client`, `TCP RST-server`, `TCP FIN`, and `Aged out`.

- PAN-DB is now the default URL filtering database.

- The default value in the WildFire Server field (**Device > Setup > WildFire > General Settings**) is changed from `default-cloud` to `wildfire-public-cloud` in order to enable support for the WF-500 WildFire appliance.

- When creating a security policy, the policy will now by default match services running on the **application-default** port for the service rather than on an any port. The **application-default** setting prevents applications from running on unusual ports and protocols, which if not intentional, can be a sign of undesired application behavior and usage. Note that when you use this option, the firewall still checks for all applications on all ports, but with this configuration, applications are only allowed on their default ports/protocols.

- The GlobalProtect gateway client **Login Lifetime** is extended to enable configuration of the validity period an IKE security association (SA) for third-party IPSec clients connecting to GlobalProtect using the extended authentication (X-Auth) capabilities. Previously, the **Login Lifetime** field only controlled the length of time a GlobalProtect agent gateway login was valid. The SA for IPSec tunnels between a third-party VPN client and the gateway was valid for eight hours before requiring an IKE phase 1 re-key, and this SA lifetime was not configurable. Now, the GlobalProtect **Login Lifetime** set on the **Network > GlobalProtect > Gateways Client Configuration** tab controls both the amount of time a gateway login is valid (for GlobalProtect agents) as well as the amount of time the IPSec client SA with the gateway is valid. In addition, to enable support for IPSec

© Palo Alto Networks, Inc.

clients that require a shorter key lifetime and to provide more configuration granularity for GlobalProtect agent configurations, the minimum value allowed for the **Login Lifetime** has been changed from 24 hours to two hours.

- If you enable the **Send Hostname in Syslog** (Logging and Reporting section of **Device > Setup > Management**) option, the default value used in the syslog header now is the fully qualified domain name (FQDN) which comprises the hostname and the domain of the device (firewall or Panorama) that sends the syslogs.

  In PAN-OS 5.0, if the **Send Hostname in Syslog** was enabled, the syslog header contained the device's FQDN when both the hostname and domain were configured or just the hostname when the domain was not configured. If neither hostname nor domain was configured, the syslog header contained the IP address of the sending device.

  In PAN-OS 4.1 and earlier versions, if the **Send Hostname in Syslog** option was enabled, the syslog header contained the IP address of the sending device.

 © Palo Alto Networks, Inc.

# Associated Software Versions

The following minimum software versions are supported with PAN-OS 6.0:

| Palo Alto Networks Software | Minimum Supported Version with PAN-OS 6.0.0 |
|---|---|
| Panorama | 6.0.0 |
| User-ID Agent | 6.0.0 |
| Terminal Services Agent | 5.0.0 |
| NetConnect | Not supported in 6.0 |
| GlobalProtect Agent | 1.2 |

© Palo Alto Networks, Inc.

# Known Issues

▲ Known Issues in PAN-OS 6.0
▲ Known Issues Specific to VM-Series Firewalls

For recent updates to known issues for a given PAN-OS release, refer to https://live.paloaltonetworks.com/t5/Articles/Critical-Issues-Addressed-in-PAN-OS-Releases/ta-p/52882.

## Known Issues in PAN-OS 6.0

| Issue Identifier | Issue Description |
|---|---|
| | While using Firefox version 30.0 with the Firebug add-on enabled, some aspects of the web interface and Panorama can be unresponsive.<br>Workaround: Disable the Firebug plug-in or downgrade your Firefox version. You can also upgrade your Firefox browser—this incompatibility is a known issue in Firefox version 30.0 and is addressed in Firefox version 31.0. |
| **68588**<br>*This issue is now resolved. See* PAN-OS 6.0.7 Addressed Issues. | If you configure a firewall as a Panorama-managed device but you do not restart the firewall after doing so, the firewall will forward predefined reports to Panorama that do not display any data.<br>Workaround: To ensure that predefined reports forwarded to Panorama are populated correctly after configuring the firewall as a managed device, restart the management server in one of two ways: either reboot the firewall or execute the `debug software restart management-server` CLI command. |
| **66059** | Regardless of the **Time Frame** you specify for a scheduled custom report on a Panorama M-100 appliance, the earliest possible start date for the report data is effectively the date when you configured the report. For example, if you configure the report on the 15th of the month and set the **Time Frame** to **Last 30 Days**, the report that Panorama generates on the 16th will include only data from the 15th onward. This issue applies only to scheduled reports; on-demand reports include all data within the specified **Time Frame**.<br>Workaround: To generate an on-demand report, click **Run Now** when you configure the custom report. |
| **63854**<br>*This issue is now resolved. See* PAN-OS 6.0.5 Addressed Issues. | For PAN-OS 6.0 release versions, virtual system administrators can perform XML API configuration commands only for the virtual systems they are an administrator for and no longer have access to XML API operational mode commands. |
| **63186** | If you perform a factory reset on a Panorama virtual appliance and configure the serial number, logging does not work until you reboot Panorama or execute the `debug software restart management-server` CLI command. |


© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **61720** | By default, the GlobalProtect app adds a route on iOS mobile devices that causes traffic to the MDM server to bypass the VPN tunnel.<br>Workaround: To configure the GlobalProtect app on iOS mobile devices to route all traffic—including traffic to the MDM server—to pass through the VPN tunnel, perform the following tasks on the firewall hosting the GlobalProtect gateway (**Network > GlobalProtect > Gateways > Client Configuration > Network Settings > Access Route**):<br>• Add `0.0.0.0/0` as an access route.<br>• Enter the IP address for the MDM server as an additional access route. |
| **59749** | On the Panorama web interface, the **Policies > Security > Post Rules > Combined Rules Preview** window does not display post rules and local rules for managed devices. |
| **58202** | On the **Monitor > Session Browser** tab, refreshing the session causes the Filter menu to be incorrectly displayed on the web interface. |
| **58049** | The Service dialog for adding or editing a service object in the web interface displays the incorrect port range for both source and destination ports: `1-65535`. The correct port range is `0-65535` and specifying port number 0 for either a source or destination port is successful. |
| **57843** | A file blocking profile configured with **continue-and-forward** as the action is not successfully sent to the WildFire cloud. |
| **56434** | When the GlobalProtect app accesses an MDM server through a Squid proxy, you must add the MDM server SSL access ports to the proxy server allow list. For example, if the SSL access port is 8443, add `acl SSL_ports port 8443` to the allow list. |
| **49742** | The following issues apply when configuring a firewall to use a hardware security module (HSM):<br>• Thales nShield Connect—The firewall requires at least four minutes to detect that an HSM has been disconnected, causing SSL functionality to be unavailable during the delay.<br>• SafeNet Luna SA—When losing connectivity to either or both HSMs in a high availability (HA) configuration, the display of information from the `show ha-status` and `show hsm info` commands is blocked for 20 seconds. |

## Known Issues Specific to VM-Series Firewalls


- VM-Series on VMware vSphere Hypervisor (ESXi)
- VM-Series on Citrix SDX
- VM-Series on VMware NSX


**VM-Series on VMware vSphere Hypervisor (ESXi)**

| Issue Identifier | Issue Description |
|---|---|
| **62573**<br>*This issue is now resolved. See* PAN-OS 6.0.3 Addressed Issues. | Upgrading a VM-Series firewall with the VM-1000-HV license can cause the firewall to go into maintenance mode without providing information about the cause.<br>The VM-1000-HV license requires 5GB RAM on the VM-Series firewall due to the higher capacities supported by this license but this memory requirement is not enforced, which makes it possible to apply the VM-1000-HV license to a VM-Series firewall without enough memory allocated (default allocated memory is 4GB). |
| **57321** | After upgrading a VM-Series firewall to PAN-OS 6.0 on an ESXi host, issuing the `request shutdown system` command does not successfully shutdown the virtual firewall; the ESXi host continues to display the virtual firewall as powered on. |

**VM-Series on Citrix SDX**

| Issue Identifier | Issue Description |
|---|---|
| **63282**<br>*This issue is now resolved. See* PAN-OS 6.0.3 Addressed Issues. | A VM-Series firewall on a Citrix SDX server will not reboot into maintenance mode to allow administrators to perform specific tasks, such as reverting images or changing between FIPS and CC mode. |
| **56452** | On the Citrix SDX server, you must disable VLAN stripping on 1-gigabit ports assigned to the VM-Series firewall if you want to configure subinterfaces with VLAN tags (divide a physical interface into multiple logical interfaces that filter traffic based on VLAN tags) on the VM-Series firewall.<br> With VLAN stripping disabled, the NetScaler VPX and the VM-Series firewall cannot share the same port.<br>10-gigabit ports work as expected and do not need to be altered. |
| **52361** | Adding or removing ports on the SDX server after deploying the VM-Series firewall can cause a configuration mismatch on the firewall. To avoid the need to reconfigure the interfaces, consider the total number of data ports that you require on the firewall and assign the relevant number of ports on the SDX server when deploying the VM-Series firewall.<br>For example, if you assign ports 1/3 and 1/4 on the SDX server as data interfaces on the VM-Series firewall, the ports are mapped to eth1 and eth2. If you then add port 1/1 or 1/2 on the SDX server, eth1 will be mapped to 1/1 or 1/2, eth2 will be mapped to 1/3 and eth3 to1/4. If ports 1/3 and 1/4 were set up as a virtual wire, this remapping will require you to reconfigure the network interfaces on the firewall. |


© Palo Alto Networks, Inc.

**VM-Series on VMware NSX**

| Issue Identifier | Issue Description |
|---|---|
| **70222** | If the password for the administrator's account on the NSX Manager contains special characters, such as `$`, Panorama cannot communicate with the NSX Manager. The inability to communicate prevents context-based information such as Dynamic Address Objects, from being available to Panorama.<br>Workaround: To fix this issue, remove special characters from the password on the NSX Manager. |
| **59856** | After deploying the VM-Series firewall, when the firewall connects to Panorama, you must issue a Panorama commit to ensure that Panorama recognizes the device as a managed device. If you reboot Panorama without committing the changes, the firewall will not connect back to Panorama; although the Device Group will display the list of devices, the device will not display in **Panorama > Managed Devices**.<br>Further, if Panorama is configured in a high availability (HA) configuration, the VM-Series firewall is not added to the passive Panorama peer until the active Panorama peer synchronizes the configuration. During this time, the passive Panorama peer will log a critical message: `vm-cfg: failed to process registration from svm device. vm-state: active`. This message is logged until you commit the changes on the active Panorama, which then initiates synchronization between the Panorama HA peers and the VM-Series firewall is added to the passive Panorama peer.<br>Workaround: To reestablish the connection to the managed devices, commit your changes to Panorama (click **Commit** and select Commit Type: **Panorama**). In case of an HA set up, the commit will initiate the synchronization of the running configuration between the Panorama peers. |
| **59573** | Live migration of the VM-Series firewall is not supported when you enable SSL decryption using the forward proxy method. Use SSL inbound inspection if you need support for live migration. |
| **58839** | When deleting the VM-Series deployment, all VMs are deleted successfully; however, sometimes a few instances still remain in the datastore.<br>Workaround: Manually delete the VM-Series firewalls from the datastore. |
| **58833** | In some scenarios, traffic from newly added guests or virtual machines is not steered to the VM-Series firewall even when the guests belong to a Security Group and are attached to a Security Policy that redirects traffic to the VM-Series firewall.<br>Workaround: Reapply the Security Policy on the NSX Manager. |
| **58832** | The VM-Series firewall fails to deploy with an error message: `Invalid OVF Format in Agent Configuration`.<br>Workaround: Use the following command to restart the ESX Agent Manager process on the vCenter Server: `/etc/init.d/vmware-vpxd tomcat-restart`. |
| **58260** | If a high availability (HA) failover occurs on Panorama at the same time that the NSX Manager is deploying the NSX edition firewall, the licensing process fails with the error: `vm-cfg: failed to process registration from svm device. vm-state: active.`<br>Workaround: Delete the unlicensed instance of the VM-Series firewall on each ESXi host and then redeploy the Palo Alto Networks next-generation firewall service from the NSX Manager. |

| Issue Identifier | Issue Description |
|---|---|
| **58170** | When the datastore is migrated for a guest, all current sessions are no longer punted to the VM-Series firewall. However, all new sessions are secured properly. |
| **58168** | When deploying the VM-Series firewall, the Task Console displays: `Error while enabling agent. Cannot complete the operation. See the event log for details`. This error displays even on a successful deployment. You can ignore the message if the VM-Series firewall is successfully deployed. |
| **57954** | If you deploy the VM-Series firewall and then assign the firewall to a template, the change is not recorded in the bootstrap file.<br>Workaround: Delete the Palo Alto Networks NGFW Service on the NSX Manager, and verify that the template is specified on **Panorama > VMware Service Manager**, register the service, and re-deploy the VM-Series firewall. |
| **57614** | When an ESXi host is rebooted or shutdown, the functional status of the guests is not updated. Because the IP address is not updated, the dynamic tags do not accurately reflect the functional state of the guests that are unavailable. |
| **57533** | The vCenter Server/vmtools displayed the IP Address for a guest incorrectly after vlan tags were added to an Ethernet port. The display did not accurately show the IP addresses associated with the tagged Ethernet port and the untagged Ethernet port. This issue was seen on some Linux OS versions such as Ubuntu. |
| **57265** | When editing a Security Policy with a network introspect rule, an `invalid (tcp) port number` error—or `invalid (udp) port number` error—displays when you remove the destination (TCP or UDP) port.<br>Workaround: Delete the network introspect rule and create a new one. |
| **57205** | When defining traffic introspection rules (to steer traffic to the VM-Series firewall) on the NSX Manager, either the source or the destination for the rule must reference the name of a Security Group; you cannot create a rule from any to any Security Group.<br>Workaround: To redirect all traffic to the VM-Series firewall, you must create a Security group that includes all the guests in the cluster. Then you can define a security policy that redirects traffic from the cluster and to the cluster. |
| **57203** | Duplicate packets are being punted to the VM-Series firewall. This issue occurs if the distributed vSwitch used for punting is enabled in promiscuous mode.<br>Workaround: Disable promiscuous mode. |
| **55586** | When adding or removing a Security Group (Container) that is bound to a Security Policy, Panorama does not get a dynamic update of the added or removed Security Group.<br>Workaround: On **Panorama > VMware Service Manager**, click the **Synchronize Dynamic Objects** link to initiate a manual synchronization to get the latest update. |
| **55393** | Dynamic Tags (update) do not reflect the actual IP address set on the guest. This issue occurs because the vCenter Server cannot accurately view the IP address of the guest. |


© Palo Alto Networks, Inc.

# PAN-OS 6.0.13 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.13 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **89962** | An update was made to ensure that firewalls and appliances running PAN-OS 6.0.13 and later PAN-OS 6.0 releases meet the transition requirements for random number generators (RNGs)—such as the X9.31 RNG—provided in SP 800-131A. The X9.31 RNG is replaced by the CTR_DRBG specified in SP 800-90A. |
| **89752** | A security-related fix was made to address a buffer overflow condition. |
| **89750** | A security-related fix was made to address a stack underflow condition. |
| **89717** | A security-related fix was made to ensure the appropriate response to special requests received through the API interface. |
| **89706** | A security-related fix was made to prevent some CLI commands from improperly executing code. |
| **77163** | Fixed an issue where the `/var/log/secure` log file inflated and consumed available disk space. With this fix, PAN-OS uses a log rotation function for this log file to avoid consuming more disk space than is necessary. |
| **67207** | Fixed an issue where an upgrade to a PAN-OS 6.0 release caused the dataplane to stop responding. |

© Palo Alto Networks, Inc.

# PAN-OS 6.0.12 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.12 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **86938** | The client certificate used by PAN-OS and Panorama to authenticate to the PAN-DB cloud service, the WildFire cloud service, and the WF-500 appliance expired on January 21, 2016. The expiration results in an outage of these services. To avoid an outage, either upgrade to content release version 550 (or a later version) or upgrade PAN-OS and Panorama instances running a PAN-OS or Panorama 6.0 release to PAN-OS (or Panorama) 6.0.12 or a later release. |
| **85721** | Fixed an issue where firewalls with a specific OCZ Deneva hard disk (model DENCSTE251M21) configured in a RAID and running PAN-OS 6.0.9 or later releases experienced RAID errors. |
| **85065** | Fixed a CLI input parsing issue that caused a process on the management plane to stop responding when processing unexpected input. |
| **83519** | A security-related fix was made to address CVE-2015-5600. |
| **81367** | A security-related fix was made to address CVE-2015-4024. |

© Palo Alto Networks, Inc.

# PAN-OS 6.0.11 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.11 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **81452** | Fixed an issue where switching context from the Panorama™ web interface to a managed firewall did not indicate whether the administrator was logged in over an encrypted SSL connection; the System log message was always `User admin logged in via Panorama from x.x.x.x using http` regardless whether the connection was encrypted. With this fix, the System log now specifically reports `User admin logged in via Panorama from x.x.x.x using http over an SSL connection` when the administrator is connected through an encrypted SSL connection to differentiate from non-encrypted connections. |
| **79443** | Fixed an issue in the web interface where, in some cases, the PHP session cookie (PHPSESSID) was not marked as secure. |
| **79367** | Fixed an issue in PAN-OS where GlobalProtect™ clients experienced delays and intermittently failed to retrieve the gateway configuration for connecting to a GlobalProtect gateway when the firewall was in a high availability (HA) configuration and under a heavy load. This issue occurred due to an issue with the synchronization of HIP reports between gateways on HA peers when there was a high number of near-simultaneous GlobalProtect connection requests. With this fix, the sync process is modified so that GlobalProtect clients are able to download the configuration and connect to the network as expected even when multiple clients are attempting to connect at the same time. |
| **79104** | Fixed a rare issue on a PA-7050 firewall where the HA1 and HA1 backup links experienced heartbeat failures that caused split brain in a high availability (HA) configuration. |
| **78652** | Fixed a rare issue where a firewall dropped URL requests when the management plane (MP) URL *trie* (data structure) reached 100% capacity. With this fix, when the MP URL trie reaches 90% capacity, URLs in the cache are cleared until the MP URL trie utilizes only 50% of capacity so that the trie cannot reach maximum capacity and cause requests to be dropped. |
| **78413** | Fixed an issue on a PA-7050 firewall with multiple virtual systems where a memory leak was observed related to the First Packet Processor (FPP) management plane process when running the `show session meter` CLI command. |
| **78304** | A security-related fix was made to address a cross-site request forgery (CSRF) issue in the web interface. |
| **77816** | Fixed an intermittent issue where some Windows 7 GlobalProtect clients using two-factor authentication (LDAP and certificate) lost connection to the portal or gateway and could not reconnect due to a failed authentication with the error `Required client certificate is not found` even when the certificate was available. |

| Issue Identifier | Description |
|---|---|
| **77548** | Fixed an issue where changing the **Configuration refresh interval** on the **Tunnel Settings** tab (**Network > GlobalProtect > Gateways > Satellite Configuration**) did not update the Refresh Time as expected. With this fix, you can click **Refresh GW Config** (**Network > IPSec Tunnels > Gateway Info**) to update the Refresh Time without having to **Reconnect to GW**. |
| **76711** | Fixed an issue where the dataplane stopped responding on a device using a shared gateway with Captive Portal in redirect mode. |
| **76615** | Fixed an issue on a PA-7050 firewall where running the `request system private-data-reset` command when there was a faulty disk drive on the Log Processing Card (LPC) caused an LPC failure during reboot. |
| **76561** | Fixed an issue where the DHCP relay agent dropped DHCPDISCOVER packets that the agent could not process due to multiple BOOTP flags. With this fix, the DHCP relay agent recognizes the first BOOTP flag in a DHCPDISCOVER packet and ignores any additional BOOTP flags that may exist (per RFC 1542) so that multiple BOOTP flags do not cause DHCPDISCOVER packets to be dropped. |
| **76083** | Fixed an issue where no System logs were generated for failed login attempts using the CLI over an SSH connection. With this fix, additional System logs now provide visibility for failed logins to the management interface even if those attempts come from a CLI over an SSH connection. |
| **75881** | Fixed an issue on a PA-5000 Series firewall where the management plane and dataplane restarted due to a race condition that occurred when the **Enforce Symmetric Return** option was enabled in the policy-based forwarding (PBF) rules (**Policies > Policy Based Forwarding > Forwarding**). This race condition caused inaccurate PBF `return-mac ager` lists, which caused the restarts. With this fix, the firewall retrieves and checks return MAC entries to avoid this race condition and associated restarts. |
| **74558** | Fixed an issue on a PA-7050 firewall where, after upgrading to a PAN-OS 6.1 release, the post-upgrade autocommit failed when the high availability (HA) peer was still running a PAN-OS 6.0 release. |
| **73878** | Fixed an intermittent issue where BGP failed to redistribute the static Discard routes as expected after a high availability (HA) failover. |
| **71674** | Fixed an issue for firewalls running PAN-OS 6.0 releases where the SSL VPN process (*sslvpn*) restarted with a core dump due to a null pointer. |
| **71611** | Introduced a CLI command in response to an issue on PA-7000 Series firewalls where logs did not always get generated or forwarded as expected when DNS response times were too slow. If you are unable to correct DNS server issues to improve response time on your network, use the `debug management-server report-namelookup` command to work around this issue by disabling DNS name lookups in reports. |
| **70919** | Fixed an issue where the dataplane in a high availability (HA) active/active configuration restarted when a `session update/remove` message was received from the peer while the session was pending an FPGA result. With this fix, FPGA results are ignored if the system receives a `session update/remove` message while waiting for those results. |
| **66681** | Resolved a dataplane restart issue due to race conditions. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **65370** | Fixed an issue where log queries utilizing DESC order or neq operator would not return expected results on M-100 or PA-7050 |
| **64266** | Fixed a rare issue where certain processes (*l3svc* and *sslvpn*) stopped responding when a Content update and FQDN refresh occurred simultaneously. |
| **62644** | When a copper SFP port was plugged in, the SFP interface's link displayed `unknown/unknown/up`; this has been updated to more accurately display `auto/auto/up`. |
| **59914** | Fixed an issue where the firewall did not remove the pan_task_x.log or .log.old files as expected when executing the `debug dataplane packet-diag clear log log` command. |

© Palo Alto Networks, Inc.

# PAN-OS 6.0.10 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.10 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **78897** | Fixed an issue where excessive ZIP traffic caused a null pointer exception in ZIP processing software, causing the dataplane to stop responding. |
| **78206** | Fixed an issue where a multi-dataplane platform did not properly free SSL proxy memory for SSL session-cache entries that included a username field that was parsed from a client certificate. With this fix, memory is freed up as expected for session-cache entries that include a username field parsed from a client certificate. |
| **77707** | Fixed an issue in PAN-OS 6.0.9 where **Threat Map** and **Traffic Map** were not appearing on the web interface under **Monitor > App Scope > Threat Map** or under **Monitor > App Scope > Traffic Map**. |
| **77237** | Fixed an issue where some PA-200 firewalls failed to boot when rebooted or power-cycled after a single PAN-OS software upgrade. You can work around this issue by performing two complete software upgrades, which corrects both u-boot images in the boot flash. With this fix, PA-200 firewalls can be upgraded to PAN-OS release 6.0.10 or higher releases and then rebooted or power-cycled normally without the need to first perform two consecutive upgrades. |
| **76238** | A security-related fix was made to address CVE-2015-1873. |
| **76099** | Fixed an issue where the dataplane restarted on a PA-7050 firewall when there was a NAT rule configured to use dynamic IP that falls back to dynamic IP and port (DIPP) NAT. |
| **75905** | Fixed an issue where a firewall failed to download the BrightCloud database via proxy after upgrading to either PAN-OS 6.0.8 or PAN-OS 6.0.9. |
| **75740** | Fixed an issue where the log-receiver crashed during a restart that happened at the same time that a NetFlow profile was removed from a security rule that was still processing traffic. |
| **74735** | Fixed an issue where a PA-7050 dataplane restarted when attempting to process jumbo frame packets (larger than 1,500 bytes). |
| **74511** | Fixed an issue where static discard routes did not get redistributed using OSPF; the routes were not injected in the OSPF link-state database (LSDB). With this fix, static discard routes are injected into the LSDB and distributed using OSPF as expected. |
| **74138** | Fixed an issue where some PA-7050 firewalls in high availability (HA) mode experienced packet buffer leaks in PAN-OS 6.0 or higher releases. The leaks occurred when interface tables, configurations, or IDs on the two HA devices lost synchronization and HA session sync messages included interface attributes that did not exist on the receiving device. Leaks also occurred when the HA pair was upgraded from PAN-OS 6.0 to PAN-OS 6.1 and interface IDs were out of sync during the time the two firewalls were not running the same software version. With this fix, packet buffer leaks caused by unsynchronized interface attributes are prevented. |

| Issue Identifier | Description |
|---|---|
| **74049** | Fixed an issue where the dataplane intermittently restarted on a PA-5000 Series firewall under heavy load conditions. This fix raises the priority of system health monitor packets so that they do not get dropped and cause the device to restart when under a heavy traffic load. |
| **73689** | Fixed an issue where traffic interruptions occurred due to nested encoding (ZIP content within chunked encoding), which caused the `SML VM vChecks` buffer pool to overflow. With this fix, new checks have been added to prevent `SML VM vChecks` buffer leaks. |
| **73605** | Fixed an issue where the User-ID process became unresponsive when trying to acquire the same lock twice with the same thread while executing the `idmgr reset` command for type `user`. |
| **73180** | Fixed an issue where, with Strip X-Forwarded-For (XFF) enabled under **Device > Setup > Content-ID**, an XFF IP address was not stripped before the packet was forwarded because the XFF header was split into two TCP segments due to an unusually long HTTP GET request. With this fix, the XFF field is stripped as expected when the header is split across two or more packets. |
| **72737** | Fixed a memory corruption issue that caused the dataplane to restart when SSL decryption was enabled. |
| **70669** | Fixed an issue where the User-ID process crashed due to bulk and incremental updates of terminal server users on the active-secondary device in a high availability (HA) active/active configuration. |
| **69837** | In response to a rare issue where a PA-200 firewall stopped processing traffic, additional troubleshooting information and some modifications to error checking and counter processes were added to help prevent this event and identify the root cause if it reoccurs. |
| **67523** | Fixed an issue where the second pair of Aggregate Ethernet (AE) interface ports did not stay down when both ports on the first AE interface went down. This issue occurred on a virtual wire (vwire) with two AE interfaces that had link-state-pass-through enabled and where both ports on one AE interface went down. With this fix, when both ports on one AE interface go down, the second AE interface ports go down and remain in powered down state until the first AE link recovers. |
| **67458** | Fixed an issue where a dataplane failed to get IP pool information from a dynamic IP and port (DIPP) source network address translation (SNAT) rule with an interface IP address. |
| **66406** | Fixed an issue where the current application version was not displayed correctly for managed firewalls when the firewall did not have a Threat Prevention subscription. |
| **66372** | Fixed an issue where some threat names did not display correctly in threat logs forwarded from the firewall when the logs were viewed on a syslog server. |
| **64887** | Fixed an issue on a PA-7050 firewall where some traffic was dropped after a configuration commit that included a change to the interface configuration. With this fix, the firewall updates current available memory as expected when changes to the interface configuration are committed. Without this fix, you can work around the issue by committing a security policy change following any commit that includes changes to the interface configuration, which prompts the firewall to update current available memory settings. |


© Palo Alto Networks, Inc.

# PAN-OS 6.0.9 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.9 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **76043** | Fixed a memory allocation issue on the PA-7050 firewall that caused intermittent connectivity for sessions inspected using SSL Forward Proxy decryption. An update was made to increase the proxy memory pool for PA-7050 firewalls, to allow for more memory to be allocated for SSL Forward Proxy sessions. |
| **74932** | Fixed an issue where high availability (HA) failovers that occurred with simultaneous route advertisements caused a routing process to restart, which then caused the firewall to restart. |
| **73813** | When using the PAN-OS® CLI in configuration mode, the `show predefined signature` CLI command incorrectly displayed App-ID™ signatures and patterns for some predefined applications. The `signature` command option has been removed and the `show predefined` command now correctly displays application information, but does not display the App-ID signature and pattern. |
| **73790** | Additional security-related enhancements were made to support frame-busting for the firewall web interface, in order to prevent framing of web interface elements. |
| **73757** | A security-related fix was made to enforce character encoding specified in HTTP headers due to CWE-116: Improper Encoding or Escaping of Output. |
| **73638** | A security-related fix was made to address issues related to HTML encoding. |
| **73337** | Fixed an issue where a VM-Series firewall with a VPN configuration restarted due to a buffer overflow caused by a race condition. |
| **73309** | Attempting to use the web interface or CLI to upload a WildFire™ content release to Panorama™ displayed an error (**Device > Dynamic Updates > WildFire**). This issue has been fixed so that WildFire content updates can be uploaded successfully to Panorama. |
| **73071** | Fixed an issue where the firewall incorrectly sent duplicate SYN packets for ftp-data sessions. |
| **72825** | Fixed an issue where traffic interruptions for various traffic patterns occurred when data was not released after packet processing. This caused vChecks to remain allocated for an extended period of time, which depleted the buffer pool. With this fix, the vCheck offset is modified so that data can be released and processed at a later time to avoid traffic interruptions. |
| **72763** | Fixed an issue where HA3 packet forwarding may fail in a high availability (HA) active/active configuration when using an Aggregate Ethernet (AE) subinterface to send and receive traffic. |
| **72737** | Fixed a memory corruption issue that caused the dataplane to restart when SSL decryption was enabled. |

| Issue Identifier | Description |
|---|---|
| **72730** | Fixed an issue where it was possible for a firewall under heavy load conditions to send malformed BGP keep-alive messages to a BGP neighbor, causing the BGP neighbor to flap. |
| **72544** | A security-related fix was made to address CVE-2014-8730. For additional information, refer to the PAN-SA-2014-0224 security advisory on the Palo Alto Networks Security Advisories web site at https://securityadvisories.paloaltonetworks.com. |
| **72536** | Fixed an issue where packet buffers leaked when a firewall with inbound SSL decryption enabled attempted to block a connection and send TCP RST packets to the connection endpoints. With this fix, TCP RST packets sent by the firewall to the connection endpoints no longer cause buffers to leak when SSL Inbound Inspection is enabled. |
| **72285** | New counters have been added to track the upload process for files forwarded to WildFire. |
| **72092** | Addressed an LSVPN issue where routes advertised by GlobalProtect™ satellites were not installed in a GlobalProtect gateway routing table. This issue has been resolved so that the GlobalProtect gateway correctly accepts routes from GlobalProtect satellites. |
| **71262** | When two M-100 appliances were in a high availability (HA) active/passive configuration, memory usage for the passive appliance increased significantly compared to the memory usage for the active appliance. This was due to a management process memory leak on the passive device and the issue is fixed. |
| **71040** | Resolved an issue that caused SFP+ ports to hang following a restart and the ports continued to stay in link down state. |
| **70903** | Fixed an issue where SNMP traps from some firewalls were not parsed correctly by the SNMP manager. |
| **70816** | Fixed an issue where the `Address is not valid` error displayed when running commands that initiate a filtering session based on an IPv6 address (examples of the affected commands include `clear session all`, `set application dump`, and `test decryption-policy-match`). IPv6 address validation now works correctly. |
| **70544** | A dataplane restart occurred when the **SSL Opt-Out** page was enabled (to notify users that SSL connections are decrypted), the RC4 cipher was enforced, and a long URL was accessed. This issue has been fixed so that the dataplane does not restart when SSL decryption is enabled. |
| **70304** | Resolved an issue that was seen with **Rematch Sessions** enabled, where a race condition could occur when new security policies were matched to existing sessions. |
| **70150** | Resolved an issue where Simple Network Management Protocol (SNMP) traps were not correctly sent to the SNMP trap destinations following a software upgrade. This issue is fixed so that SNMP traps are generated and correctly sent to SNMP trap destinations after performing an upgrade. |
| **69900** | Fixed an issue where the tech support file did not contain some expected files, including /var/log files. |
| **69737** | On platforms with multiple dataplanes, stale IPv6 neighbor entries were not removed and replaced with new IPv6 neighbor entries when the IPv6 neighbor table was full. This issue has been fixed so that stale IPv6 neighbor entries are correctly removed when the neighbor table threshold is reached. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **69242** | When a user failed to authenticate using the web interface, firewall system logs did not display the user's source IP address. Updates have been made so that a failed authentication on the web interface is logged with two entries. One entry is logged as a `general` event and displays only the username of the user who failed authentication. The other entry is logged as an `auth-fail` event and displays both the username and source IP address of the user who failed authentication. |
| **68982** | Fixed an issue where the firewall stops receiving new reports from WildFire when the report ID on the WildFire public cloud exceeds a certain limit. (WildFire continues to generate reports but does not log them as expected on the firewall). |
| **67930** | Fixed an issue where an update to a stale IPv6 neighbor entry could cause a dataplane restart. |
| **67810** | When a PA-5000 Series device initiates sessions on different data planes in an environment with multiple virtual systems, session traffic sometimes failed to span across virtual systems. This issue has been resolved so that inter-virtual system sessions succeed with a dynamic network address translation (NAT) policy configuration. |
| **66217** | Fixed an issue that caused unreliable VoIP communication. The issue occurred for firewalls in a high availability (HA) active/active configuration, and was due to predict session traffic not being installed correctly on the HA peer when the peer was setup to use a dataplane other than DP0 to process the session. |
| **64759** | Fixed an issue where a high availability (HA) failover occurred due to insufficient kernel memory on a PA-5000 Series firewall that was attempting to handle unusually heavy network and system traffic. With this fix, the kernel memory on PA-5000 Series firewalls is increased to ensure sufficient kernel memory is avail-able for ping requests and keep-alive messages even when under an unusually heavy load. |
| **58547** | Policy-based forwarding (PBF) with symmetric return did not work when the traffic was translated with source NAT. Return traffic, which needs to be forwarded using the same interface on which it arrived, was dropped with the message `Symmetric Return: Packet dropped, no return MAC found`. The issue is fixed. |

© Palo Alto Networks, Inc.

© Palo Alto Networks, Inc.

# PAN-OS 6.0.8 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.8 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **73111** | Dataplane restarts were caused by a race condition between dataplane packet processes, where the session resource allocation became out of sync between central processing units (CPUs). A fix was added to keep session resource allocation in sync between dataplane processes. |
| **72241** | Following an upgrade, attempting to perform a high availability (HA) configuration sync between two HA peers in an active/passive or active/active deployment did not work correctly. This issue has been fixed so that HA peers will sync correctly following an upgrade. |
| **72068** | If a firewall with Open Shortest Path First (OSPF) enabled was then restarted, a flapping condition was seen between the firewall and the adjacent OSPF neighbor, and a new OSPF election was forced by the firewall. This issue has been fixed so that following a firewall restart, any OSPF adjacency remains established. |
| **71939** | Addressed an issue where enabling a second Network Processing Card (NPC) on a PA-7050 firewall resulted in URL packets being dropped by the second NPC and URL lookups could fail. This issue has been fixed so that URL lookups are performed correctly and web pages load quickly. |
| **71850** | Changing the IP address for a log card interface on a PA-7050 firewall caused an issue where traffic log forwarded to syslog servers stopped until the firewall was restarted. This was due to an issue where the firewall sent out traffic using an internal IP address (which was recognized as an invalid source IP address by devices intermediate to the firewall and the syslog server) following a change to the log card interface IP address. This issue has been fixed so that changing the IP address for a log card interface does not cause the firewall to send out traffic using an internal IP address. |
| **71688** | On a PA-7050 firewall with OSPF enabled, a restart caused OSPF neighbor adjacency states to flap. This issue was caused by an incorrect slot number setting on the Network Processing Card (NPC) for the session owner. With this fix, the NPC slot number for the session owner is properly selected and OSPF neighbor adjacency is established. |
| **71553** | Fixed an issue where dataplane processes restarted when handling SSL decryption sessions during high availability (HA) message updates. The fix for this issue included the addition of a global counter. |
| **71512** | A fix was made to add frame-busting to the firewall web interface to prevent framing of web interface elements. |
| **71503** | Addressed an incorrect file permissions issue in the web interface. |
| **71486** | A fix was made to address an issue with user input sanitization to prevent Cross-site Scripting (XSS) attacks against the web interface. |

| Issue Identifier | Description |
|---|---|
| **71464** | If the firewall initiates a point-to-point protocol over Ethernet (PPPoE) session, an issue was seen when a server responds with a PPPoE Active Discovery Offer (PADO) packet that was greater in size than the maximum transmission unit (MTU) of the firewall interface. In this case, the PADO packet was dropped. This issue has been addressed so that PADO packets are handled correctly by the firewall, including when the size of the packet is greater than the MTU for the firewall interface. |
| **71333** | In a high availability (HA) active/active configuration with an IPSec tunnel configured to terminate on a floating IP address, Encapsulating Security Payload (ESP) was performed by the device that did not own the floating IP address. The encapsulated packets failed the IPSec anti-replay check on the remote end of the IPSec tunnel and were discarded. With this fix, packets are always sent to the owner of the floating IP address to be encapsulated. |
| **71321** | Removed support for SSL 3.0 from the GlobalProtect™ gateway, GlobalProtect portal, and Captive Portal due to CVE-2014-3566 (POODLE). |
| **71320** | Removed support for SSL 3.0 from the web interface due to CVE-2014-3566 (POODLE). |
| **71273** | A security-related fix was made in PAN-OS to address issues related to parsing XML data. |
| **71199** | In a Large Scale VPN (LSVPN) setup, a GlobalProtect satellite reconnecting to a GlobalProtect gateway after receiving a different IP address, changed the GlobalProtect routing metrics when installing the gateway access routes into the satellite routing table. With this fix, the original gateway routing priority is restored when the GlobalProtect satellite reconnects to the GlobalProtect gateway with a different IP address. |
| **71148** | When attempting to add an address to an **Address Group** using the Panorama web interface, filtering for the address returned no results even though the address object did exist and was displayed as configured on the **Objects > Addresses** page. Additionally, filtering for the same address object when attempting to add the address to a security rule displayed different results for the address object name. This issue has been resolved so that filtering for an address correctly displays any configured addressed objects, and address object names are displayed consistently. |
| **70820** | Addressed an issue for PA-7050 firewalls, where Real-time Transport Protocol (RTP) predict sessions remained in the Opened session state and did not become an active session. This caused the RTP packets to not merge correctly with the predict session and the packets were dropped if they did not specifically match to an allow policy. |
| **70499** | Fixed an issue where traffic that matched to a predict session and then converted to a flow session was then being incorrectly matched to security policies where the only matching criteria defined in the policy was a custom application. A fix was made to perform a second policy lookup after predict session traffic is converted to flow session traffic. |
| **70383** | When using the Panorama™ XML API to register an IP address to a Dynamic Address Group on a targeted firewall, an error was displayed that the user was not authorized to perform the operation. This issue has been resolved so that using the XML API to register an IP address to a Dynamic Address Group on the firewall results in the firewall correctly registering the IP address and updating the membership information for the dynamic address group. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **69934** | Fixed an issue where active File Transfer Protocol (FTP) enabled with Dynamic IP Network Address Translation (NAT) failed. Inconsistent translation was seen for PORT packets sent over the control connection established between the FTP server and the NAT device. |
| **69685** | Updates were made to existing Russian time zones and new Russian time zones were added to the available list of global time zones for a device, to accommodate the 2014 changes to Russian time zones. |
| **69528** | An issue was seen with captive portal enabled in an environment with multiple virtual systems, where two virtual systems were configured as User-ID collectors. When captive portal timed out a User-ID entry in one virtual system, captive portal did not time out for the same User-Id entry in the second virtual system, and the user was not prompted to re-authenticate to the captive portal login page. This issue has been fixed so that if captive portal times out for a User-ID entry in one virtual system, the same User-ID entry correctly times out in the second virtual system and the user is prompted to authenticate to the captive portal login page. |
| **68764** | When a DNS proxy server is configured on the firewall, the proxy settings were not used and DNS resolution was requested to resolve service.brightcloud.com. After the fix, the connection request by the firewall to BrightCloud is always forwarded to the proxy. |
| **68702** | An error was displayed when pushing a policy from Panorama to a managed firewall with a user group defined in the policy. The error displayed was Duplicate group name and this issue has been resolved so that pushing a user group from Panorama to a managed firewall works correctly. |
| **68560** | Addressed an issue where vulnerabilities were logged as unknown when an ampersand character (&) was used in the **Comment** field when creating a custom vulnerability object (**Objects > Custom Objects > Vulnerability**). Using the ampersand character in the **Comment** field when creating a custom vulnerability object is supported, and does not cause the vulnerability to display as unknown. |
| **68430** | The dataplane restarted unexpectedly due to a lack of memory. An update has been made to provide additional debug information for this issue. |
| **68372** | Setting up a static MAC configuration for a tagged interface configured on a VLAN did not work correctly. This occurred when a process that communicates between the dataplane and the management plane restarted, and the issue has been resolved. |
| **68371** | Addressed an issue where you could not install the BrightCloud database when the default url-db was set to PAN url-DB, and you had not downloaded the BrightCloud database previously. |
| **68100** | The Strip X-Forward-For (XFF) feature did not correctly remove the contents of the XFF header field when the HTTP GET request contained encoded URL characters. This issue has been fixed to ensure that the Strip XFF feature removes the contents of the XFF field from the HTTP header, including when the HTTP GET request contains encoded URLs. |
| **67885** | Panorama predefined reports for vulnerabilities were inconsistent with the predefined report for Vulnerabilities on the managed firewall. This issue has been addressed so that reports are correctly synchronized between Panorama and managed devices. |

© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **67861** | Following an upgrade to a PAN-OS 6.0 release version, in some cases virtual wire interfaces went down after restarting the firewall. This issue has been fixed so that the status for virtual wire interfaces is no longer down after upgrading to a PAN-OS 6.0 release version and restarting the firewall. |
| **67567** | When a new version of the BrightCloud URL database was downloaded and installed, if the category for a URL changed in the new version of the database from the previous version, the change in category was not reflected on the dataplane. With this fix, URL categories on the dataplane are updated correctly after a installing a new version of the BrightCloud database. |
| **65477** | A deadlock issue has been addressed where an inter-process communication (IPC) failed while another IPC operation was in progress. This resulted in the management server restarting with a core dump. |
| **64958** | Resolved an issue displayed on the Panorama web interface, where the **Oldest Log** statistics incorrectly showed no data (**Panorama > Managed Collectors > Collector Statistics**). |
| **64930** | Dynamic objects could be lost if the device server restarted unexpectedly. This has been fixed so that dynamic objects are repopulated if the device server process unexpectedly restarts. |
| **64600** | When a dynamic block list was configured on the firewall to be updated according to a list hosted on a web server, the firewall was configured to access the web server using a proxy server. An issue was seen where the firewall was connecting directly to the web server hosting the block list, instead of using the proxy server to connect to the web server. This issue has been fixed so that when a proxy server is configured, the firewall updates the dynamic block list using the proxy server and does not connect to the server hosting the block list directly. |
| **64040** | The logging partition could become full on a log collector if the total configured quota was larger than 90%. A fix was made to address the way free space is allocated for the logging partition. |
| **62791** | An update was made to reduce the number of TCP stale sessions for PA-5000 Series devices. |
| **61201** | Scheduled email reports were not being delivered, though the reports were generating and displaying correctly on the **Monitor** tab on the web interface. This issue was due to a memory leak for a back-end process that maintains configuration information for the firewall. This issue has been fixed so that scheduled email reports are correctly delivered to email. |
| **58055** | Addressed an issue where a race condition between a high availability (HA) synchronization and a web interface commit caused a management server restart. |
| **55249** | You can now run the `test <feature>` CLI command for the following features: botnet, cp-policy-match, custom-url, data-filtering, decryption-policy-match, dns-proxy, dos-policy-match, global-protect-mdm, global-protect-satellite, nat-policy-match, nd, pbf-policy-match, pppoe, qos-policy-match, routing, scp-server-connection, security-policy-match, stats-service, tag-filter, url, url-info-cloud, url-info-host, user-id, vpn, wildfire. |


© Palo Alto Networks, Inc.

# PAN-OS 6.0.7 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.7 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **70588** | Fixed an issue where a browser with Transport Layer Security (TLS) 1.2 enforced could not access the GlobalProtect™ portal login page. This issue occurred when no client certificate was configured. |
| **70165** | Fixed an issue for PA-7050 firewalls in a high availability (HA) active/active configuration, where IPv6 fragments could cause a Network Processing Card (NPC) to restart. |
| **69956** | Resolved an issue where NetFlow information for some sessions was not being forwarded for PA-5000 Series devices due to a session ID format change. |
| **69035** | When using the **ACC** tab on the Panorama™ web interface to view statistics for a custom application, using applications filters (such as Category, Subcategory, and Technology) to filter the data displayed for the custom application resulted in no data being displayed. This occurred when Panorama was selected as the **Data Source** for the traffic data displayed on the **ACC** tab, and the issue has been resolved. |
| **68836** | In a a high availability (HA) configuration, a path monitoring failure lead to a delayed HA failover. An update has been made to optimize HA failover time. |
| **68588** | Predefined reports for a firewall connected to Panorama were not being displayed correctly if the management server for the firewall was not restarted after connecting to Panorama. This issue has been fixed so that predefined reports on a managed firewall are still displayed correctly after establishing a connection with Panorama. |
| **68472** | Addressed an issue where some expected counters were not returned in the output for the `<show> <interface>` XML API command for loopback, VLAN, and tunnel interfaces. |
| **68389** | The Application subcategory is listed as unknown in the PDF report for custom applications pushed from Panorama. This issue was resolved by correcting the report daemon to properly parse the configuration objects pushed from Panorama. |
| **68055** | Mac clients were incorrectly unable to access certain websites that Windows clients were able to access. This issue occurred when fragmented traffic passed through the firewall and the first fragment did not include the header; this caused packets to be dropped. The issue has been resolved. |
| **67864** | When a rule pushed from Panorama is selected on a managed device, the clone button in a security policy is enabled; however, rules pushed to a managed device from Panorama cannot be cloned on a managed device. With this fix, the clone button for rules pushed from Panorama correctly shows as disabled on the web interface for a managed device. |
| **67856** | When a primary NTP server was down and the secondary NTP server was up, the CLI incorrectly showed the secondary NTP server as down. This has been fixed so that the CLI displays the correct status for both the primary and secondary NTP servers. |

| Issue Identifier | Description |
|---|---|
| **67833** | When a tech support file was generated for Panorama, private information was not being removed correctly from files within a device group if the device group had a space in its name. With this fix, device groups with a space in the device group name are handled correctly when generating a tech support file. |
| **67814** | Panorama displayed the secure-proxy-password in the web interface under **Panorama > Setup > Services** and in the CLI. With this fix, Panorama encrypts the secure-proxy-password and downgrade attempts to versions that show the secure-proxy-password will fail until you remove the secure-proxy-password from the configuration. |
| **67788** | The configuration log on Panorama displayed the secure-proxy-password. With this fix, the configuration log encrypts the secure-proxy-password. |
| **67778** | Fixed an issue where a GlobalProtect Mac client experienced intermittent connectivity with the GlobalProtect portal when running Mac OS X 10.10. |
| **67455** | Made an update to the enforcement for the SSL Inbound Inspection setting `block when resources are unavailable` so that hosts cannot resume an SSL session, when that session has been removed from the SSL-decrypt session cache due to the cache being full. The host must start a new session to continue. |
| **67300** | Addressed an issue on the VM-Series firewalls where enabling packet capture for certain application-level gateway (ALG) traffic causes the system to restart. |
| **67268** | A DNS sinkhole was not responding to DNS queries that contained DNSSEC Header flags, such as the Authenticated Data (AD) bit or Checking Disabled (CD) bit. The DNS sinkhole now correctly responds to DNS queries that contain DNSSEC header flags (AD or CD bit) and the firewall can successfully provide the DNS sinkhole IP address to the client. |
| **67258** | The process (*mprelay*) that communicates between the dataplane and the management plane restarted unexpectedly. A policy-based forwarding (PBF) rule configured with symmetric return but not specifying an IPv6 next hop address resulted in excessive neighbor discovery (ND) update messages that caused a conditional loop, which then caused the mprelay process to restart. With this fix, the IPv6 ND performs correctly and avoids unexpected restarts of the mprelay process even if no IPv6 next hop address is specified. |
| **67187** | The following error was displayed due to an issue that caused a User-ID process to restart: `Abnormal system memory usage detected, restarting userid with virtual memory`. Many GlobalProtect users logging into the system, and the resulting high availability (HA) synchronization of the HIP reports, caused the virtual memory to exceed its limit. |
| **67160** | An issue was fixed where a data port configured on a PA-7050 firewall as a Log Card Interface was restarting when commits were performed. |
| **66924** | When logging in to the Panorama web interface with two-factor RADIUS authentication, Panorama would successfully authenticate an administrator but then immediately log the administrator out of the web interface. With this fix, Panorama no longer logs the administrator out of the web interface following a successful authentication. |
| **66918** | Memory corruption issues related to SSL decryption caused the dataplane to restart and resulted in a flapping condition between firewalls in an HA cluster. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **66635** | Enabling SSL Forward Proxy decryption with a self-signed certificate could sometimes cause the certificate presented to the client to have a negative serial number, causing an error on the client. |
| **66482** | In some cases, the web interface was not accessible for an M-100 appliance even though the appliance could be accessed through the CLI. This issue is the result of a race condition, where a reporting process attempted to delete session that had been previously deleted. This issue is not addressed so that the web interface and CLI are both accessible for an M-100 appliance. |
| **66364** | Fixed an issue that prevented two certificates with the same subject name from being installed following an upgrade to a PAN-OS 6.0 release version. |
| **66208** | A brute-force attack on an unprotected management interface on the firewall caused the /var/log/btmp log file to inflate and consume available disk space. With this fix, PAN-OS enables a log rotation function for failed SSH logins, such as those from brute-force attacks. |
| **66161** | For a PA-3050 firewall, the next hop MAC address for hardware offload did not update correctly for routes that were learned using multiple dynamic routing protocols. This occurred when the routing protocols were not using the default administrative distance. Packets used the wrong egress interface when the session was offloaded. This issue has been resolved. |
| **65859** | Fixed an issue where the dataplane could restart when SSL Forward Proxy decryption was enabled and a certain packet sequence was received. |
| **65850** | Addressed an issue where a high availability (HA) backup failed due to there being no buffer space available. |
| **65180** | Removed support for the RC4-MD5 and RC4-SHA cipher suites. These ciphers cannot be used to negotiate an SSL/TLS connection to the GlobalProtect portal. |
| **64713** | Removed support for the RC4-MD5 and RC4-SHA cipher suites. These ciphers cannot be used to negotiate an SSL/TLS connection to the management interface of the firewall.. |
| **64606** | When navigating to the GlobalProtect portal using a browser that had Transport Layer Security (TLS) 1.2 enabled, and when using a client certificate for authentication, the SSL connection failed due to issues with fallback to a lower TLS version. With this fix, fallback to a lower TLS version succeeds with Google Chrome and Mozilla Firefox. A specific behavior of Internet Explorer still exhibits issues. |
| **64223** | Fixed an issue where FQDN objects that were added to a dynamic address group were not listed after issuing the `request system fqdn show` command, with the command displaying a message that no FQDN object is used in the policies. |
| **63150** | In an Active/Active HA setup, User Datagram Protocol (UDP) sessions with a certain traffic pattern caused the session state to flap frequently and generate excessive traffic logs. This issue is now fixed and the session state is stable. |
| **61056** | IPSec tunnels could not be established due to the firewall running out of encap/decap context. An enhancement was made to allow more granular debugging for the depletion of encap/decap context for IPSec tunnels, in order to provide further root cause analysis. Restart the firewall as a workaround for this issue. |

© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **60617** | TCP connectivity issues occurred on a virtual machine when SYN Flood was enabled with the SYN Cookies **Activate Rate** set to zero. This was due to an issue where active and passive FTP were not working under different NAT configurations, and has been addressed so that both active and passive FTP work correctly under different NAT configurations and do not cause TCP connectivity issues. A workaround for this issue is to set the **Activate Rate** to a value greater than zero (**Objects > Security Profiles > DoS Protection > Flood Protection > SYN Flood**). |


© Palo Alto Networks, Inc.

# PAN-OS 6.0.6 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.6 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **68899** | Fixed an issue that affected PA-7050 firewalls. An HSCI port configured as an HA2 interface went down due to a dataplane board restarting. An improvement has been made so that if more than one dataplane board is installed, an HA2 interface on an HSCI port will stay up. |
| **68708** | Addressed the bash vulnerability CVE-2014-7169 that relates to how environment variables are processed when the shell starts up. This fix prevents a user with an account on the firewall, from using the vulnerability to gain escalated privileges. |
| **68286** | An issue was seen where setting up a password for a proxy server caused the management plane to restart (**Device > Setup > Services > Proxy Server**). This was due to a process that was restarting when the password was configured. This issue has been fixed. |
| **67910** | While viewing the **Combined Rules Preview** on the Panorama™ web interface (**Policies > Preview Rules > Combined Rules Preview**), the **Device** drop-down displayed the virtual system number for some virtual systems and not the virtual system name. This issue has been fixed so that the **Device** drop-down in the Combined Rules Preview displayed the names of the virtual systems that you can select. |
| **67873** | Resolved an issue where disabling OSPF and performing a commit caused a routing process to restart. This was due to a race condition related to running the `show route` command during a commit. |
| **67860** | On the Panorama web interface, the **Preview Changes** button could not be clicked when **Group HA Peers** was selected. This has been fixed so that the **Preview Changes** button is correctly enabled when **Group HA Peers** is selected (**Panorama > Templates**). |
| **67723** | Fixed an issue that occurred when an OSPF profile was configured with a tunnel interface and the **Passive** state was disabled for the tunnel interface. In this case, the runtime configuration incorrectly overrode the configured settings, and the tunnel interface was incorrectly enabled as **Passive**. |
| **67516** | Fixed an issue with a high availability (HA) active/active configuration where a physical MAC address was returned for a floating IP address instead of a virtual MAC address. This has been addressed so that the floating IP correctly responds to ARP requests with a virtual MAC address. |
| **67483** | This fixes an issue where a firewall failed to email scheduled reports due to a race condition. |
| **67436** | The `debug software trace reportd` and `debug software core reportd` commands were added to the CLI command structure. |
| **67399** | Fixed an issue that occurred in a high availability (HA) active/active configuration, where log card interfaces were synced to the HA peer and resulted in duplicate IP addresses. |

| Issue Identifier | Description |
|---|---|
| **67344** | Fixed an issue for the M-100 appliance where the `show log-collector detail` command was presenting incorrect information. |
| **67069** | Internal packet path monitoring failure errors caused the dataplane to restart. An enhancement was made to detection and recovery mechanisms to minimize the impacts of these errors. |
| **66959** | Addressed an issue where configuring overlapping IPv6 addresses when adding a **OSPFv3 range** on the **Network > Virtual Routers > OSPFv3 > Area > Ranges** tab caused a restart to occur after attempting to commit. This issue has been fixed so that if an invalid range is entered, an error is displayed; if a valid range is entered, you can continue to commit successfully. |
| **66690** | The dataplane restarted due to a process restart while running traffic in a high availability (HA) configuration. Additional checks have been added to avoid a possible race condition, which had led to the dataplane process restart. |
| **66503** | Addressed an issue where the firewall dataplane experienced an out of memory condition, which could cause the dataplane to restart and the firewall to go into nonfunctional state. |
| **66466** | Addressed an issue for PA-2000 Series firewalls, where a device failed to handle a volume of packets (larger than the MTU) on the management interface. Symptoms of this issue included device unresponsiveness, a random restart, traffic failures or ATA errors on the console. This issue has been resolved. |
| **66220** | An issue was seen in a high availability (HA) active/passive configuration where the secondary device was not able to pass traffic after a failover until a routing process was restarted. This issue has been fixed so that when a failover occurs, the secondary device correctly becomes the Backup Designated Router (BDR). |
| **66168** | Resolved an issue where adding an FQDN to the Servers table to specify an LDAP server in an LDAP server profile caused intermittent connection issues (**Device > Server Profiles > LDAP**). |
| **66025** | Configuration files with names longer than 32 characters were allowed and could be successfully imported, but load and delete operations would fail. With this fix, configuration file names of up to 32 characters can be imported and configuration files with longer names are prevented from being imported with an error. |
| **65607** | If the last remaining user was removed from the **Allow List** for an LDAP authentication profile (meaning no users remained on the list), authd was not notified that the group was empty and retained the last user's information. That user could continue to be authenticated despite no longer being included in the allow list. This has been addressed so when the last user remaining on the allow list is removed, the user can no longer authenticate. |
| **65565** | Fixed an issue where selecting R**eplay attack detection** in the **GlobalProtect gateway satellite** configuration on the web interface did not actually enable replay attack detection. |
| **65488** | Resolved an issue that occurred with a high availability (HA) active/passive configuration, where using the `request -availability state suspend` command to suspend the active peer and perform an HA failover resulted in some packet loss. |
| **65176** | Resolved an issue that caused a dataplane restart on the VM-Series firewall, when using RC4 ciphers to decrypt SSL traffic. This issue is specific to ESXi servers using AMD processors. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **64647** | Following an upgrade from a PAN-OS 5.0 release version to a PAN-OS 6.0 release version, high availability (HA) synchronizations were failing intermittently. The HA sync failures were due to an issue where the host information profile (HIP) database was not syncing correctly during the HA sync between the peers. This issue has been fixed so that performing a HA synchronization works correctly. |
| **64379** | This fixes an issue where older cached IP address to username mappings on a User-ID agent could overwrite newer IP address to username mappings on the firewall. With this fix, IP address to username mappings with more recent timestamps take precedence over IP address to username mappings with older timestamps. |
| **64309** | A WildFire™ threat log with the severity `Informational` was incorrectly showing the severity `Emergency` when forwarded to a syslog server. This issue has been fixed so that when WildFire logs are forwarded to a syslog server, the log entries show the correct severity for the log. |
| **62768** | Unreliable DNS servers incorrectly provide NXDOMAIN responses. To help prevent incorrect WildFire sample categorization, NXDOMAIN responses are no longer shared across WildFire samples. Each NXDOMAIN response will be evaluated on a sample by sample basis. |
| **61205** | Using the web interface to export traffic logs in CSV format was showing an error that the query job failed. This issue has been addressed so that exporting traffic logs to CSV works correctly. |
| **58820** | On PA-5000 Series firewalls, Static Source NAT, Dynamic IP NAT, and Destination NAT session processing has been enhanced to greatly improve the throughput in these NAT scenarios, allowing the firewall to use multiple CPUs to process NAT sessions, rather than anchoring the sessions to a CPU based on destination IP hash. |

© Palo Alto Networks, Inc.

© Palo Alto Networks, Inc.

# PAN-OS 6.0.5-h3 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.5-h3 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

Before upgrading to PAN-OS 6.0.5-h3 or subsequent PAN-OS 6.0 release versions, and if you have asymmetric routes in your network or have attached a zone protection profile, first refer to the Important notes in the Upgrade to PAN-OS 6.0 section of the PAN-OS New Features Guide Version 6.0.

| Issue Identifier | Description |
|---|---|
| **69173** | Under certain conditions, unspecified layering of packet-level evasions could be used to bypass signature matching of the session. |

© Palo Alto Networks, Inc.

# PAN-OS 6.0.5 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.5 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **68018** | When configured in high availability (HA) mode, the High Speed Chassis Interconnect (HSCI) ports using QSFP modules on the PA-7050 experienced packet loss. This HSCI link flapping issue has been resolved. |
| **67613** | Resolved an issue where uploading files to Wildfire sometimes failed when the files were uploaded from the client to the server using HTTP. |
| **67565** | For Panorama and PAN-OS, when an Admin Role Profile was set up with Privacy permissions disabled, the administrator was able to search the traffic logs based on username. This issue has been fixed so that an administrator without privacy permissions cannot search traffic logs based on username. |
| **67378** | Resolved an issue where LEDs for copper ports on the PA-7050 firewall that were configured as high availability (HA) ports were not correctly showing active links. |
| **67246** | Resolved an issue where modifying an existing policy caused any tags configured in that policy to be removed. |
| **67073** | Fixed an issue where link flapping was seen during a high availability (HA) failover. |
| **67058** | Resolved an issue where a Vulnerability Protection profile did not match a custom vulnerability signature when the vulnerability name was entered as the **Threat Name** in the vulnerability protection rule. This issue has been resolved so that the Vulnerability Protection profile matches signature when the vulnerability name is entered as the Threat Name in the rule. As a workaround, the **Threat Name** in the Vulnerability Protection profile could be set to **any** so that the profile will match the custom signature. |
| **66910** | With Block Password Change Period enabled to specify an amount of time during which an administrator cannot change a password, an administrator's password could still not be changed after the period had passed. This has been resolved so that when the period has passed, a password for an administrator can be successfully modified. |
| **66692** | Resolved an issue where device administrators with privacy settings disabled were still able to view usernames on the **ACC** tab. |
| **66610** | Scheduled email reports for Panorama were not correctly showing IP addresses resolved to hostnames, and were continuing to display IP addresses for entries. This issue has been resolved so that hostnames are correctly displayed in scheduled email reports instead of the corresponding IP address. |

| Issue Identifier | Description |
|---|---|
| **66342** | Resolved an issue that occurred with Panorama peers in a high availability (HA) configuration. Following an HA failover, a managed device forwarded logs to the secondary Panorama while the primary Panorama was down. When the primary Panorama came back online, the managed device incorrectly continued to forward logs to the secondary Panorama. This has been fixed to ensure that managed devices correctly forward logs to the primary Panorama when both the primary and secondary Panoramas are online. |
| **66215** | Resolved an issue that caused a dataplane restart when a tech support file was being generated and a content update was being performed at the same time. |
| **66199** | When the option **Target to all but these selected devices** was enabled for a Panorama Security policy, the policy was being pushed to the dataplane of managed devices that were explicitly excluded from receiving the policy (the policy should not have been pushed to these devices at all). This issue has been fixed so that when **Target to all but these selected devices** is enabled, the policy is correctly pushed to devices associated with that policy, and is not pushed at all to devices that are explicitly excluded. |
| **66043** | A PA-5050 firewall was showing that it only supported 32,768 device routes when it should support a maximum of 64,000 routes. This issue has been resolved to ensure that the firewall supports 64,000 device routes. |
| **65936** | Addressed an issue where the device server could crash if a configuration referenced different names for a single GlobalProtect gateway entry. |
| **65905** | Addressed an issue where a routing process restarted, causing the firewall to restart. This occurred when configuring virtual routers with both OSPF and BGP enabled with aggregate routes and routes to be redistributed. |
| **65833** | In a high availability (HA) active/active configuration, resolved an issue where the primary HA device did not resume the primary status after being suspended, when the preemptive option was enabled on both firewall peers in the setup. |
| **65828** | On Panorama, with the option to **Share Unused Address and Service Objects with Devices** enabled, certain conditions caused unused objects to still be pushed to managed devices. This issue has been fixed so that unused objects are not pushed to firewalls when the option to **Share Unused Address and Service Objects with Devices** is not selected. |
| **65801** | Traffic for a voice and video application was being dropped because Network Address Translation (NAT) was not correctly translating the application's IP address and the traffic was not able to be routed correctly. This issue has been resolved so that H225 and H245 traffic is handled efficiently by the firewall-NAT is correctly performed for voice and video applications and calls placed using Voice over IP (VoIP) do not fail. |
| **65765** | Addressed an issue related to parsing long URLs which in some cases was causing an infinite loop that could cause the dataplane to restart. |
| **65747** | When one port configured as a virtual wire interface was brought down and then brought back up, both virtual wire interfaces continued to flap before stabilizing. This was an issue for 10G Fiber ports and has been resolved. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **65678** | After upgrading to Panorama 6.0.0 or a later release version, previously authenticated non-local administrator accounts, like RADIUS authenticated administrators, could not log in using the CLI unless first logging in to the web interface. This issue has been addressed so that, following a Panorama upgrade, an authenticated administrator can successfully log in to the CLI without first logging in to the web interface. |
| **65511** | The Mozilla compatibility user agent headers in Internet Explorer 11 were changed, causing Captive Portal to identify Mozilla as an unsupported browser and NTLM authentication was not performed. This issue has been resolved so that Internet Explorer 11 can perform NTLM authentication. |
| **65496** | The Panorama command line interface (CLI) was displaying an extra line when exiting the CLI when running Panorama 6.0.2 or 6.0.3. The additional line, which indicated that the management session was terminated, was not displayed in releases prior to Panorama 6.0.2 and has been removed. |
| **65340** | A commit failed without a clear error message when a static route statement overlapped an OSPF export statement. Updates have been made to clarify the error message showing conflicts that could lead to a failed commit. |
| **65292** | Safe Search Enforcement was not being enforced when users were logged into either Bing or Google accounts (if users were logged out of their Bing or Google accounts, Safe Search enforcement worked correctly). This issue has been resolved so that Safe Search Enforcement blocks Bing or Google search results correctly, even when users are logged into Google or Bing accounts. |
| **65256** | The system logs on a PA-7050 displayed an error informing the user that an executable file could not be found. This issue has been fixed so that the executable file has a valid file path and the error is no longer displayed. |
| **65068** | Resolved an issue seen while forwarding logs to a syslog server, where the `risk-of-app` field was displayed as an application instead of a risk value. |
| **64990** | Attempting to sort entries for scheduled custom reports according to the **Receive Time** under **Monitor > Reports** was not the correct sorted results. This has been fixed so that these results can be correctly displayed in ascending or descending order by **Receive Time**. |
| **64908** | With both OSPF and BGP enabled, if OSPF was configured to redistribute the default route `0.0.0.0/0`, this caused OSPF to redistribute routes that were configured for BGP route redistribution. Resolved this issue so that OSPF no longer redistributes routes configured for BGP redistribution. |
| **64707** | Fixed an issue where a device server process restarted unexpectedly while processing the URL cache. |
| **64705** | Exporting logs from Panorama and then attempting to import them to a managed device did not display the logs on the managed device. Imported logs were placed in a temporary directory before being added to the main database. An ordering error in this process was cleaning up the temporary directory before the imported logs were correctly merged into the main database. This issue has been resolved. |
| **64664** | Resolved an issue where the Detailed Log View for traffic logs was not displaying the Log Links widget (**Monitor > Logs > Traffic**). |

© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **64642** | Fixed an issue where an IP address obtained from a DHCP server on a dataplane interface was not binding to the dataplane and an IP address obtained from a PPPoE server on a dataplane interface was not bound to the dataplane. This issue has been resolved. |
| **64596** | A user icon was showing for a group object, when the group object was added to a security policy. This has been fixed so that a group icon is shown for a group object when added to a security policy. |
| **64550** | Resolved an issue where enabling NAT64 resulted in unexpected ARP behavior on Layer 3 interfaces. |
| **64165** | Following a Panorama upgrade to PAN-OS 6.0.2, several duplicated email notifications were being sent for old logs. This was due to an issue with sequence and acknowledgment (SEQ/ACK) numbers. Updates have been made to reduce the duplicated logs. |
| **64076** | When a template with the L3 Forwarding Enabled option selected was created on Panorama running a Panorama 5.1 release version, commits could not be performed after then upgrading to the Panorama to run PAN-OS 6.0.2. This was due to a migration issue that occurred during the upgrade and was related to the L3 Forwarding object. This issue is now resolved. |
| **63854** | Resolved an issue where a virtual systems administrator could not use the XML API to add IP address to username mappings. |
| **63395** | Custom reports configured on Panorama were not being pushed to a PA-7050, and were displaying on the PA-7050 with no data. This has issue has been resolved so that custom reports configured on Panorama are successfully pushed to a PA-7050 and display correctly. |
| **62565** | Resolved an issue where OSPFv3 was not redistributing a default route correctly. |
| **62375** | The GoDaddy root certificate authority (CA) was missing from the list of trusted certificate authorities. When SSL decryption was configured, sites using the GoDaddy root certificate authority were displayed as not trusted. The GoDaddy Root Certificate Authority - G2 has been added to the list of trusted CAs. |
| **61038** | Resolved an issue where the device server would sometimes fail to send a URL request to BrightCloud, and the URL categories would display as not resolved. |
| **60030** | URL Filtering logs now display when a search has been successfully blocked due to Safe Search Enforcement. An entry with a `block-url` action is listed in the URL Filtering logs whenever a user is presented with a block page (**Monitor > Logs > URL Filtering**). |


© Palo Alto Networks, Inc.

# PAN-OS 6.0.4 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.4 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **66157** | Fixed an issue on a PA-7050 firewall, where traffic through an IPSec tunnel failed during tunnel key renegotiation due to a race condition. |
| **65643** | When a firewall was configured with both 1GB and 10GB interfaces, and a 10GB interface went down, a back-end process was not reporting the link down event instantly. An update was made to prevent back-end processes from delaying to report an interface's link status as down. |
| **65612** | Fixed an issue that occurred when Panorama was running a PAN-OS 6.0 release version and a managed firewall was running a PAN-OS 5.0 release version. Reports generated for WildFire private cloud were not accessible on Panorama and showed the error message: `the requested resource does not exist`. |
| **65557** | In certain situations, a firewall in a high availability (HA) configuration could exit a suspended state without user intervention. With this fix, additional checks have been implemented to ensure that a firewall can exit a suspended state only due to an administrator request using the CLI or web interface. |
| **65534** | Resolved an issue that caused a PA-7050 to restart due to a link failure. This was due to an issue where when a link failure occurred, only one side of the link was being reset. Link detection has been updated so that if a link failure occurs, both sides of the link are reset. |
| **65445** | Addressed an issue where the ACC charts on the firewall web interface did not display any data after the order of the columns was modified and the width of the columns was adjusted. |
| **65396** | Fixed an issue that occurred when the option to **Share Unused Address and Service Objects with Devices** was disabled on Panorama. When this option was disabled, attempting to commit a Device Group to a managed firewall failed. |
| **65328** | Addressed an issue with Panorama where logs were not being uploaded. The use of a space character in Log Collector Group names led to this issue. You can no longer include a space character in a Log Collector Group name. |
| **65316** | The firewall web interface would allow you to add an **Exempt IP Address** to a vulnerability from the Threat Logs but did not save the settings and enforce the IP exception unless an **Exempt Profile** was also selected. The web interface has been updated so that you cannot add an **Exempt IP Address** without also selecting an **Exempt Profile**. |
| **65278** | Selecting a service object displayed in the Service column for a Security policy rule and then selecting **Value** did not expand to display the service object's members as expected (**Policies > Security**). This has been updated so that selecting a service object in the Service column correctly displays the service object's **Value** fields. |

| Issue Identifier | Description |
|---|---|
| **65156** | A security-related fix was made to address CVE-2014-0224. For additional information, refer to PAN-SA-2014-0003 on the Security Advisories site (https://securityadvisories.paloaltonetworks.com). |
| **65146** | Resolved an issue that occurred in a high availability (HA) active/active configuration, where HA3 packets could drop for a few seconds on the active-primary device while a content update was running on the active-secondary device. |
| **65098** | Resolved an issue that was seen after a shared policy was pushed from Panorama, where all traffic through a firewall stopped. This was due to an issue where the routing loop configuration between two Virtual Routers was causing incorrect programming to the offload processor on the PA-4000 Series firewall. |
| **64960** | In a high availability (HA) active/active configuration, when a device received Jumbo packets that had been fragmented, it transmitted the packets as fully assembled Jumbo packets. This issue has been resolved so that fragmented Jumbo packets are no longer transmitted as assembled Jumbo packets; both fragmented and assembled packets are transmitted correctly. |
| **64870** | Resolved an issue that was seen in a Panorama high availability (HA) active/passive configuration, where saving the configuration on the active device did not correctly sync the predefined configuration to the passive device. |
| **64795** | Custom application port settings were getting deleted inadvertently by clicking the **OK** button after opening the specific custom application. This issue has been addressed so that clicking the **OK** button after opening a custom application entry does not delete port settings. |
| **64656** | Addressed an issue that occurred on PA-3000 Series firewalls, where the dedicated HA2 port was taking more than 10 seconds to link up when an Ethernet cable was connected. |
| **64591** | The fix in this bug resolves two issues:<br>• PA-3000 Series firewalls could stop passing traffic due to errors in the internal data path. A registry value was changed in the dataplane packet processor to prevent this error.<br>• When highly compressed traffic is passing through PA-7050, PA-5000 Series, or PA-3000 Series firewalls, this caused some additional latency in other traffic passing through the firewall. This fix prevents the latency from being introduced to the traffic by changing the amount of data being handled by the signature matching FPGA. |
| **64553** | When using the PAN-OS integrated User-ID agent, after the mapping has timed out, the firewall continued to send WMI probes to devices that were no longer responding. With this fix, the firewall does not perform WMI probing after the time-to-live (TTL) timer for an IP address to username mapping expires. |
| **64438** | Static DNS proxy entries configured on the DNS Proxy page are set with an initial TTL value of 7 days. For static DNS proxy entries, the TTL counter should always remain at the initial value; however, in this case, the TTL counter for the static DNS proxy entries was counting down to zero. This issue has been resolved so that the TTL value for static DNS proxy entries remains at 7 days. |
| **64322** | Resolved an issue where Panorama was running a PAN-OS 6.0 release version and could not push Dynamic Address Groups to a firewall running a PAN-OS 5.0 release version. This issue has been resolved and was due to migration issues caused by the removal of the Dynamic Address Object functionality in PAN-OS 6.0. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **64289** | Addressed an issue where SSL Certificate Common Names containing uppercase characters were not parsed properly due to case sensitivity. In certain conditions, this caused the URL Categorization lookup to fail and the issue has been resolved. |
| **64166** | After approximately 388 days of uptime, the firewall lost the IP address to username mappings on the dataplane. This issue has been addressed so that the firewall does not lose IP address to username mappings when it reaches this uptime. |
| **64122** | Addressed an issue where a User-ID process unexpectedly stopped on the passive device in a high availability (HA) active/passive configuration. This was due to an issue caused by a restart of the Windows server that was hosting the Terminal Server Agent. |
| **63971** | Addressed an issue where the Network File System (NFS) process failed intermittently on PA-5000 Series devices configured in a high availability (HA) configuration. This issue occurred because HA session updates were not being processed in the correct order and sessions were stalled. Though the issue was not specific to NFS traffic, it was more prevalent for NFS traffic and other application traffic with long session timeout periods. |
| **63658** | An update was made to support the export of threat packet capture (PCAP) files using the XML API for a role-based user. |
| **63620** | Fixed an issue with PA-3000 Series firewalls where traffic could stop passing through the firewall or the dataplane could restart due to errors seen on the field-programmable gate array (FPGA). |
| **63423** | Resolved an issue with the firewall web interface where attempting to filter traffic logs using the parameter `user in` caused the management plane to stop responding. |
| **63304** | A dataplane process stopped unexpectedly when processing malformed inbound-proxy SSL encrypted handshake packets. Defensive protection logic has been incorporated to gracefully handle these packets and to avoid process failures. |
| **63113** | Resolved an issue where MIB walk counters were showing different information than the output of the `show session info` command on the firewall. |
| **63110** | Changed the level of a User-ID log to show as a debug log instead of a warning log. |
| **62741** | The feature that allows you to configure polymorphic objects in the Panorama web interface was not working in PAN-OS 6.0 release versions (the feature worked as expected in PAN-OS 5.0 release versions). This feature allows you to configure objects that do not exist in policies on Panorama if the following command is configured: `debug skip-policy-address-check yes`. This issue has been fixed so that the polymorphic objects feature works correctly. |
| **62410** | With SSL decryption enabled on the firewall, certain web sites that use Transport Layer Security (TLS) Renegotiation were failing to load because the sites' certificates included renegotiation extensions. An update was made so that the firewall supports TLS Session Renegotiation. |
| **62323** | Made fixes to improve the issue where the firewall went into a non-functional state due to an out of memory condition caused by an internal process. Updates have been made to resolve some of the memory utilization issues. |

© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **62305** | In a multiple virtual system environment, subsequent commit failures were seen following an upgrade to a PAN-OS 6.0 release version. This occurred when prior to the upgrade, Panorama had pushed security policies that reference tags. This issue has been resolved so that following an upgrade to PAN-OS 6.0.4 or later release version commits are successful and tag objects pushed from Panorama prior to the upgrade are automatically created. |
| **62244** | When trying to add an address object to a policy, the search drop-down failed to display the configured address object. This issue was only seen when the number of address objects available exceeded 500 objects, and has been addressed so that address objects are displayed correctly when using the search drop-down to add the object to a policy. |
| **61972** | When highly compressed traffic passes through a PA-7050, PA-5000 Series, or PA-3000 Series firewall, this caused some additional latency in other traffic passing through the firewall. This fix prevents the latency from being introduced to the traffic by changing the amount of data being handled by the signature matching FPGA. |
| **61846** | The firewall web interface required the field **Peer HA1 IP Address** to be populated in order to disable high availability (HA). This has been addressed so that the field **Peer HA1 IP Address** is not mandatory to disable HA. |
| **61170** | The PAN-OS 6.0 SNMP MIB has been updated to add 42 new system log traps. The new traps are added as follows:<br>• 5 traps added under PAN_ELOG_EVENT_HA<br>• 3 traps added under PAN_ELOG_EVENT_VPN<br>• 2 traps added under PAN_ELOG_EVENT_USERID<br>• 11 traps added for PAN_ELOG_EVENT_MDM (new category)<br>• 21 traps added for PAN_ELOG_EVENT_RAID (new category)<br>• 2 traps added for PAN_ELOG_EVENT_VM (new category)<br>You can download the updated SNMP MIB (PAN-MIB-MODULES-6.0-RevC.zip) from the following location: https://live.paloaltonetworks.com/docs/DOC-6587. |
| **61641** | Addressed an issue where incorrect entries and information were displayed when moving between pages in the **Monitor** tab on the web interface. |
| **61632** | Addressed an issue where a service route was configured for a syslog server using a destination IP address, but the syslog server's traffic was using the management port. |
| **61329** | Using the web interface to add an IP address to a security policy, where a subnet was added without a proper subnet mask (for example, `1.1.1.1/`), was not blocked as an error. The expected behavior would be to block adding an IP address ending with a `/` as that character could translate to a value of any on the dataplane. An update was made to check for the `/` character. |
| **59795** | For administrators that logged into the web interface using RADIUS authentication, configuration changes did not cause the **Commit** button to be lighted (the **Commit** button showed as disabled) and were not displayed on the **Device > Config Audit** page. This issue did not exist for administrators authenticating using the local database. This issue has been resolved for all authentication methods (local user database, RADIUS, and LDAP). |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **59715** | On PA-5000 Series devices, when sessions were initiated on different dataplanes in an environment with multiple virtual systems, sometimes session traffic failed to span across virtual systems. This issue has been resolved so that inter-virtual system sessions succeed with a destination NAT (DNAT) policy setup. |
| **57687** | An update was made to enhance initialization of the PAN-DB cloud lookup process during device startup. |
| **56786** | Security rules with a URL category attribute were not being correctly matched against certain sessions with `set deviceconfig setting ssl-decrypt url-proxy` enabled, which sets a proxy for SSL sessions if the URL category is blocked. An update has been made so that a security policy lookup is performed after the URL category is received for proxy purposes, in order to ensure the information is retrieved correctly. |
| **40883** | Addressed a root partition that was at 100% usage due to an issue where FireMon processes (where FireMon was being used as an external third-party monitoring utility to the Palo Alto Networks firewall) continued to run on a device and did not close correctly, even after the FireMon connection had timed out. |

© Palo Alto Networks, Inc.

© Palo Alto Networks, Inc.

# PAN-OS 6.0.3 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.3 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **63862** | A back-end process was using an excessive amount of memory, causing an out of memory condition. An update has been made to improve how out of memory conditions are addressed. When an out of memory condition occurs, the back-end process will be terminated in place of any other critical processes. This ensures that any critical processes will continue to run despite an out of memory condition, and the back-end process will later restart automatically. |
| **63859** | A PA-500 device was not sending out scheduled email reports. This was due to a rare issue where report generation was suspended and a race condition caused it to be unable to resume, resulting in subsequent scheduled reports failing to generate. This issue has been addressed so that the report generation process is able to correctly resume after being suspended. |
| **63635** | Addressed an issue where Panorama restarted unexpectedly due to a back-end reporting process that stopped responding. |
| **63626** | The clear log operational command was not available on the PA-7050 firewall in the first three releases of PAN-OS 6.0 (6.0.0, 6.0.1, and 6.0.2). As of PAN-OS 6.0.3, this command is now available. |
| **63608** | Upgrading from a PAN-OS 5.0 release version to a PAN-OS 6.0 release version failed when the PAN-OS 5.0 configuration included a custom application configuration that included the spyware identification option (CLI command: `spyware-ident`). This was related to the spyware identification option not being supported in PAN-OS 6.0 release versions until PAN-OS 6.0.3. This issue has been addressed so that upgrading to PAN-OS 6.0.3 or higher release versions is successful, even when the PAN-OS 5.0 configuration includes a custom application configuration with the spyware identification option. |
| **63602** | An issue was resolved where a User-ID process on the firewall stopped working when a connected User-ID Agent sent an XML API logout message containing a valid IP address but no entry name. |
| **63587** | Custom signatures added as exceptions to a Vulnerability Protection Profile were not being displayed on the **Exceptions** tab on the Panorama web interface (**Objects > Vulnerability Protection > Exceptions**). This issue has been addressed so that custom signatures added as exceptions to a Vulnerability Protection profile are correctly displayed on the **Exceptions** tab. |
| **63582** | On the Panorama web interface, attempting to filter target devices to which specific security rules have been applied did not display results. This issue has been addressed so that the filter works correctly. |
| **63570** | Addressed a race condition that caused scheduled report generation to stall preventing scheduled reports from generating. |

| Issue Identifier | Description |
|---|---|
| 63551 | Selecting **Sync to Peer** on the **High Availability** widget on the web interface **Dashboard** did not display the expected pop-up to confirm the sync to peer. This issue and another minor display issue were seen using an Internet Explorer 9 or Internet Explorer 10 browser to use the web interface. These display issues have been addressed. |
| 63472 | Smart card authentication for GlobalProtect was displaying the following error message: `Authentication failed: empty password`. This issue has been addressed so that authentication is successful and the password field is correctly populated. |
| 63422 | Scheduled reports were not showing resolved host names. This was caused by an attribute value that was incorrectly set in one of the functions responsible for resolving the hostname for a given IP address. The value for this attribute has been corrected. |
| 63383 | Addressed an issue where DNS Proxy did not use the expected DNS server (according to the configured DNS Proxy rule) to perform a name server lookup. |
| 63316 | Addressed an issue that occurred when attempting to log in to the Panorama web interface. After authenticating login credentials and displaying the message: `Creating administrative session. Please wait…`, a white screen was displayed instead of the Panorama web interface **Dashboard**. This issue has been resolved so that logging in to the Panorama web interface correctly displays the **Dashboard**. |
| 63298 | Addressed an issue where syncing a configuration between Panorama high availability (HA) peers did not work correctly. |
| 63282 | Fixed an issue where a VM-Series firewall on a Citrix SDX server would not reboot into maintenance mode to allow administrators to perform specific tasks, such as reverting images or changing between FIPS and CC mode. |
| 63247 | During Session Initiation Protocol (SIP) registration, if there was no port specified in the SIP payload, the device did not add any NAT information to the SIP payload. This caused some SIP servers to use more aggressive registration intervals using the RTP Control Protocol (RTCP). This has been fixed by adding a new application: `sipviaheader-nat`. An Application Override policy is required to use this new application. |
| 63242 | A field-programmable gate array (FPGA) change was made to handle microbursts of multicast traffic and improve recovery mechanisms in latency conditions. |
| 63161 | Resolved an issue where the NetBIOS name was removed from the IP address to username mapping when the NetBIOS name contained a period punctuation mark. |
| 63147 | Fixed an issue in the migration script for the Routing Information Protocol (RIP) export rules from PAN-OS 5.0 release versions to PAN-OS 6.0 release versions that caused Panorama template commits to fail. |
| 63106 | Addressed an issue that affected an upgrade from a PAN-OS 5.0 release version to a PAN-OS 6.0 release version, where an IPSec validation error occurred when the dynamic peer type on an IKE gateway had the Local and Peer IDs defined. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **63086** | Resolved an issue that occurred in a high availability (HA) active/active configuration, where the parent application sessions on a HA peer were not updated with traffic from the child application session and the parent applications sessions timed out. This caused the parent applications to close on both devices in the HA pair. This has been updated so that when the parent application session is refreshed on the primary device, it is now also refreshed on the secondary device. |
| **63052** | When attempting to connect to GlobalProtect with token-based authentication configured, and a token could not be found, different error messages were previously displayed depending on the GlobalProtect or PAN-OS release version. This error message has been updated to be consistent across release versions. |
| **62995** | Addressed an issue where the Application usage and Traffic summary by URL Category information in User Activity Reports was not in sync with the information displayed on the **ACC** tab in the web interface, in the **Top Applications** and **URL Filtering** sections. |
| **62985** | An issue with the FPGA code for some PA-3000 Series devices was causing a small amount of VLAN tagged and offloaded packets to be truncated by 4 bytes. |
| **62969** | Online Certificate Status Protocol (OCSP) requests were using the OCSP location in the certificate instead of the location configured on the firewall. This has been adjusted so that the OCSP configuration on the firewall will take precedence. If no OCSP location is configured on the firewall, the certificate OCSP location will be used. |
| **62895** | Addressed an issue where members of groups on non-Active Directory LDAP servers were not displayed as group members on the firewall. |
| **62883** | Addressed an issue where fiber port down detection on SFP ports on a PA-2000 Series device took 10 - 20 seconds. |
| **62876** | A Panorama role-based administrator with privileges for Device Groups and Templates could view devices which were not included in any device groups or templates. An update has been made in the web interface to eliminate visibility into devices that are not a part of a device group for administrators limited to Device Groups and Templates privileges. |
| **62827** | Addressed an upgrade schema issue for Panorama running PAN-OS 6.0.1 where address objects were causing commit errors when attempting to push a configuration from Panorama to managed devices. |
| **62801** | Traffic was blocked during SSL decryption when a certificate was signed with an untrusted CA and when the option to use OCSP to check certificate status was enabled (**Device > Setup > Session > Certificate Revocation Checking**). This has been resolved so that decryption works with untrusted certificates, with the option to use OCSP to check certificate status enabled. |
| **62763** | Group mapping queries failed if the group name, or Active Directory organizational unit (OU) path, contained characters that were used in the standard LDAP query syntax (for example, parenthesis). These characters were not escaped correctly during the query building process causing the query to fail. An update was made so that the special characters are escaped. |
| **62663** | Resolved an issue where the status for Tor was inaccurately reported in the outputs for the `show wildfire status` and `test wildfire tor` CLI commands. |

© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **62596** | A shadow rule warning message was incorrectly displayed and was not accurate in indicating that a security rule's match criteria was met by a preceding rule (the match criteria had not been met by a preceding rule). The incorrect error messages were being displayed while performing a commit and a fix was made to eliminate the incorrect warning messages. |
| **62595** | Addressed an issue that was seen when using a filter to define search criteria for WildFire logs on the web interface. When adding a log filter for WildFire logs, selecting Category as the Attribute to include in the search displayed URL categories to select from in the **Value** column (**Monitor > Wildfire > Add Log Filter**). The **Value** column should have only displayed the options Malicious or Benign to select from. |
| **62591** | During extremely high rates of logging over an extended period of time, traffic and threat logs could not be displayed on the web interface or the CLI. This was due to the dataplane logging process consuming too much memory when attempting to write very large amounts of traffic logs and has been resolved. |
| **62586** | When upgrading Panorama, commits to device groups would fail if Panorama had Service Groups and Address Groups configured with the same name. This has been resolved so that group objects are independent and unique names are not required. |
| **62573** | Fixed an issue where upgrading a VM-Series firewall with the VM-1000-HV license caused a firewall to go into maintenance mode without any indication of why it did so. This issue occurred because the VM-1000-HV license requires 5GB RAM to support higher capacities but, prior to PAN-OS 6.0.3, the memory requirement was not enforced, making it possible to apply the VM-1000-HV license to a VM-Series firewall without enough memory allocated (default allocated memory is 4GB). With this fix, a memory requirement check is added so that the VM-Series firewall will boot to maintenance mode with an error informing you that you need to increase the allocated memory on the firewall to 5GB or more. |
| **62559** | Resolved an issue where creating or modifying policy rules took a long time for large configurations (for example, a configuration with 3000 rules and 20,000 objects). An update was made to increase the speed for creating and modifying policy rules for large configurations. |
| **62552** | With both SSL Inbound Inspection and virus scanning enabled, the firewall was able to detect a virus being sent, but did not block the virus from being forwarded to the target. With this fix, the virus file is not forwarded to the target. |
| **62545** | An update was made to support the export of threat packet capture (PCAP) files using the XML API. |
| **62391** | When displaying the Detailed Log View of a traffic log entry, the value of the attribute **Receive Time** was incorrect and did not match the value from the **Receive Time** column on the traffic log list. With this fix, the value of the **Receive Time** attribute in the Detailed Log View is shown correctly. |
| **62377** | Following an upgrade to PAN-OS 6.0.1, the following unexpected system log was generated daily: `Got batch event for unknown child acc_rollup`. An update has been made so that the log is no longer generated daily. |
| **62366** | The dataplane restarted with multiple cores due to a race condition caused by updating the aged out session. This issue has been fixed so that the race condition is resolved and will no longer trigger a restart. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **62315** | Resolved an issue for devices running PAN-OS 5.0 or PAN-OS 6.0 release versions, where the disk quota displayed on the Panorama web interface for a Log Collector group was not reflective of the actual disk quota output. |
| **62291** | An issue was occurring where any type of commit triggered the `syslog-ng` configuration to be reloaded and the connection to be reset. This issue was fixed so that only commits related to the syslog server configurations and syslog-related log forwarding configurations will trigger a `syslog-ng` configuration reload and the connection to be reset. |
| **62219** | Unexpected XML data was present when loading an existing configuration version from the device. This resulted in a commit failure that cited the unexpected text. An update has been made to avoid this condition and to ensure a successful commit when loading an existing configuration version. |
| **62179** | In certain situations, a tunnel interface appeared to be up and was present in the forwarding table, although the underlying IPSec tunnel was down. This issue occurred when the IPSec tunnel had a VPN Monitor profile configured with Fail Over enabled, and the device was restarted while the remote IPSec endpoint was down. With this fix, the status of the tunnel interface matches the state of its corresponding IPSec tunnel when a VPN Monitor profile with Fail Over enabled is configured, and regardless of the initial state of the tunnel. |
| **62113** | Resolved an issue for a PA-7050 device, where TAP Mode traffic did not load balance properly on a Network Processing Card (NPC), and caused a much higher load on one dataplane in comparison to another dataplane. |
| **61987** | Resolved an issue where virtual firewalls running PAN-OS 6.0 release versions did not correctly sinkhole suspicious DNS queries. |
| **61953** | When the IPv6 address `::/0` was configured under the Permitted IP Addresses section in an Interface Management profile, IPv6 communication to the interface was not allowed. Since IPv6 address `::/0` is equivalent to `any`, IPv6 communication should be allowed from any IPv6 host. This has been addressed so that you can add the IPv6 address `::/0` under Permitted IP Addresses for an Interface Management Profile to successfully allow IPv6 communication from any IPv6 host. |
| **61879** | IPSec VPN traffic could not be initiated following a restart of the `sysd` process. This issue has been resolved so that IPSec VPN traffic can be initiated even if `sysd` process is restarted. |
| **61864** | Fixed an issue where a PA-5000 Series firewall was dropping Skinny Call Control Protocol (SCCP) traffic after an upgrade to PAN-OS 6.0.0; SCCP sessions were timing out within 5 seconds due to communication issues between dataplanes. With this fix, the parent session timeout value is set appropriately and SCCP traffic is not dropped. |
| **61855** | Resolved an issue where multiple configuration pushes from Panorama to a managed device could cause that managed device's management server to restart. When this occurred, a system log would be indicating that the management server was restarted due to high memory utilization. This issue is resolved so that commits from Panorama to managed devices succeed for multiple configuration pushes and the management server remains stable. |
| **61827** | Addressed an issue where a virtual wire with **Link State Pass Through** enabled had one interface that remained connected despite the other interface being physically disconnected. |

© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **61785** | Resolved an issue where an administrator for a single virtual system could delete logs for other virtual systems using the CLI. |
| **61781** | Resolved an issue that occurred when using the Panorama web interface to set up SNMPv3 with a view. Adding a **View** with a space in the view's name caused SNMP polling to fail. This has been resolved so that an error is displayed if a space is used when naming a **View** and a new name can be entered without a space (**Panorama > Setup > Operations > Miscellaneous > SNMP Setup > Views**). |
| **61757** | The `show arp management` command timed out after 30 seconds without returning output if both primary and secondary DNS servers configured on the firewall were unreachable. The timeout has been increased up to 120 seconds to return ARP entries with or without DNS resolved names. The time taken to return the output could vary based on the number of DNS servers configured and the availability of the DNS servers to resolve host names. |
| **61724** | When a custom report was exported to CSV format, the device name in the report was showing the device's serial number instead of the device name. This has been fixed so that exporting a custom report to CSV format correctly shows the device name in the device column. |
| **61643** | Resolved an issue where email alerts were displaying all values on a single line. This occurred when the email alerts were based on custom log formats that were configured with a `\n` or a new line as a separator between fields. |
| **61350** | Following an import of a license key file from the CLI, several checks were performed to verify the certificates validity and an error was seen where the file size of the key file could be misinterpreted and the file incorrectly rejected. This issue has been resolved so that importing a valid license key file is successful. |
| **61329** | Using the web interface to add a an IP address to a security policy, where a subnet was added without a proper subnet mask (for example, `1.1.1.1/`), was not blocked as an error. The expected behavior would be to block adding an IP address ending with a `/` as that character could translate to a value of any on the dataplane. An update was made to check for the `/` character. |
| **61326** | Addressed an issue that occurred in a high availability (HA) configuration, where link monitoring for a 10G port displayed the wrong status for the port when the port was down. |
| **61284** | Traffic using policy-based forwarding (PBF) with symmetric return enforced experienced packet loss due to failed return MAC address lookups. This issue was caused by another change made in PAN-OS 5.0.11 which changed the PBF return MAC learning from the forwarding state to the session start; the PAN-OS 5.0.11 change that caused this issue has been reversed, which fixes the issue. |
| **61276** | Resolved an issue that occurred while replacing a device following an RMA. After issuing the `replace device` CLI command, exporting the device state did not show the appropriate configuration fields to set up the replacement device. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **61254** | H.323 video calls were getting dropped intermittently when video call traffic was creating a multi-level application-level gateway (ALG) session dependency pattern. The issue was observed when the grandparent session had a low timeout value and then timed out, causing its dependent parent and grandchild sessions to be terminated and the video call was dropped. This has been addressed so that TTL counters of both parent and grandparent sessions refresh and do not time out while the corresponding grandchild session continues to see traffic. |
| **61101** | Resolved an issue where the Address Resolution Protocol (ARP) table was not updated following Spanning Tree Protocol (STP) re-convergence in the network. |
| **61058** | Resolved an issue that occurred when a NAT Policy rule was configured to use a Dynamic IP pool for IP NAT translations and a fall back pool was also configured to perform **Fall Back Dynamic IP Translation** if the Dynamic IP pool ran out of addresses. In certain cases, the fall back pool was being used for translations when the IP pool was not exhausted. |
| **60781** | A Dynamic Block List address object (0.0.0.0) was being added to security rules where that Dynamic Block List was referenced, even though that object did not exist in the Dynamic Block List source file. This was due to issues with EBL Refresh, and has been addressed so that when a Dynamic Block List is referenced in a security policy, the Dynamic Block List objects displayed in the policy are accurate. |
| **60634** | Inter-vsys multicast traffic caused the dataplane to restart. The restart was due to a fragmented packet that was encapsulated into a Protocol Independent Multicast (PIM) packet and routed to another virtual system. The fix prevents this kind of fragmentation of the packets. |
| **59852** | Fixed a high dataplane CPU usage issue seen on PA-5000 Series devices that was caused by FPGA lock-up. |
| **59331** | Syncing a high availability (HA) configuration failed. This occurred when the MTU size on the management interface was set to a value smaller than the default MTU size, and dataplane interfaces were using the MTU size value set on the management interface. This resulted in dataplane interfaces (in this case, including the HA interfaces) to drop packets that were larger than the MTU size set on the management interface. |

© Palo Alto Networks, Inc.

© Palo Alto Networks, Inc.

# PAN-OS 6.0.2 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.2 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Description |
|---|---|
| **62578** | A Panorama administrator with privileges to manage Device Groups and Templates was unable to perform management and administrative functions due to several buttons and functionalities appearing to be disabled on the web interface. The issue has been resolved so that a Device Groups and Templates admin can successfully perform management and administrative functions as defined by the admin role. |
| **62567** | An issue was addressed where commits were failing intermittently following an upgrade to PAN-OS 6.0.1. This was due to the management server failing to send the ID request to the client device and has been fixed so that commits are successful following an upgrade. A second issue was addressed with this fix, where high availability (HA) synchronization was showing an error message when there was no URL Filtering license; this occurred when the PAN-DB cache was empty and the passive device treated the empty files as corrupt files. |
| **62386** | The URL category match within a security policy was not enforced, allowing non-HTTP traffic to match the rule. This update ensures proper enforcement of the URL category selection in a security policy. |
| **62352** | Addressed an issue where templates pushed from Panorama were missing BrightCloud to PAN-DB conversion files, resulting in errors when managed devices attempted to translate URL filtering policies pushed from Panorama. |
| **62321** | The Panorama web interface was inaccessible due to a corruption in MongoDB journal files. A fix was made so that if a MongoDB journal file becomes corrupt, Panorama now recovers correctly by removing the journal files and then recreating them. |
| **62244** | When trying to add an address object to a policy, the search field failed to return the address object. This issue was only seen when the number of address objects available exceeded 500 objects. |
| **62218** | Fixed an issue where an administrator with full privileges on Panorama could not delete most of the configuration on the **Network** tab of the Panorama web interface. This issue has been addressed so that all objects on the **Network** tab can be created, modified, and deleted. |
| **62208** | Resolved an issue where service routes for email alerts and SNMP messages did not take precedence if the management interface was disconnected. |
| **62199** | Panorama failed to commit when attempting to push a configuration to a managed device. This occurred when the pushed configuration included a decryption profile; even though there were no decryption port mirroring options selected in the decryption profile, the managed device did not support decryption port mirroring and the commit failed. This was due to Panorama not pruning out the portion of the configuration that was specific to decryption port mirroring and has been fixed to ensure that configuration parameters are removed when pushed to devices that do not support them. |

| Issue Identifier | Description |
|---|---|
| **62084** | Addressed an issue where domain names greater than 63 characters were getting truncated in the resolv.conf file for Linux and Mac OS even when separated by commas. |
| **62076** | After creating an Anti-Spyware profile with the **Action on DNS queries** set to **Block** (**Objects > Security Profiles > Anti-Spyware > Anti-Spyware Profile > DNS Signatures**), the threat logs for a suspicious DNS query displayed the action taken on the threat traffic as `reset-both`. This has been updated so that when the action is set to **Block** for DNS queries in an Anti-Spyware profile, the threat logs will show the action taken on the suspicious DNS traffic as `drop-all-packets` when using UDP (the threat logs will still show the action taken on suspicious DNS traffic as `reset-both` when using TCP). |
| **62057** | FTP log export failed when an explicit destination path was defined. In cases where multiple CSV files were used, the first CSV file was created but subsequent CSV files failed. This was updated to allow the creation of multiple CSV files during file export. |
| **62020** | When a device is configured to use a Time Zone with a name that is abbreviated to more than three characters (for example, America/Anchorage abbreviated to AKST), the web interface displayed an error message on the Dynamic Updates page. This issue did not affect already configured automatic updates. Workarounds for this issue include using the CLI to configure and view dynamic updates and re-setting the Time Zone of the device to a Time Zone with a name that is three characters or less, for example, America/Los Angeles abbreviated to PST. This issue has been addressed so that the Time Zone and the corresponding Time Zone name the device is configured to use do not cause an error on the Dynamic Updates page. |
| **61985** | The GlobalProtect portal became inaccessible using the web interface or the GlobalProtect client after an interface address change. This normally occurred when an interface was DHCP enabled and an address renewal occurred. The issue was caused by an address map change which published incorrect interface data to the SSL VPN component. The address map change operation has been corrected to publish the correct data. |
| **61951** | SSL Inbound Inspection sometimes failed due to SSL session reuse. This issue has been resolved so that SSL Inbound Inspection is successfully enforced according to the configured decryption policy. |
| **61992** | PA-4000 Series firewalls running PAN-OS 6.0 in a high availability (HA) active/active configurations utilizing Virtual or Floating IP addresses experienced intermittent connectivity. (HA active/passive configurations on PA-4000 Series firewalls running PAN-OS 6.0 releases were not affected by this issue.)<br>Workaround (limited): Configure static ARP entries that are bound to the physical interfaces; however, this creates the possibility of losing failover capabilities in HA active/active configurations and force traffic to prefer one firewall over another. |
| **61897** | The following error message displayed when loading Enterprise 6.0 SNMP MIBs using the ManageEngine SNMP MIB Browser: `Could not parse the file PAN-COMMON-MIB. Couldn't resolve these sequence constructs`. This issue has been resolved by fixing capitalization errors on the MIBs. |
| **61844** | Under extremely heavy load with fully Layer 7 traffic, the PA-7050 firewall was experiencing packet loss due to the depletion of the software buffer. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **61840** | Auth and priv passwords for SNMPV3 setup on managed collector groups were displayed in clear text. With this fix, the passwords are masked and are not shown in clear text. |
| **61828** | TCP SYN cookies were activated sooner than the **Activate Rate** set in the DoS protection profile. Using DoS Protection with SYN cookies now activates correctly according to the configured SYN count. |
| **61747** | The use of a 64-character string as a RADIUS shared secret was not supported. A fix is implemented to ensure a higher value for the buffer to store and encode shared secrets. |
| **61742** | With safe search enabled, when a user was attempting to access a regional site (for example, www.google.co.uk), a block page appeared directing the user to update the preferences page for the main site and not the regional site (for example, the block page directed the user to google.com preferences instead of www.google.co.uk). However, updating the preferences page for the main site did not enforce the preferences on the regional site. This has been fixed so that block pages display the correct preferences page for regional sites. |
| **61714** | Traffic logs were not displayed on a PA-7050 firewall despite traffic being generated. This was due to a disk failure and there was no system log was generated for the disk failure. This has been addressed so that a system log is generated to notify if disk failure event occurs during raid rebuilding. |
| **61696** | When using SSL Forward Proxy decryption with self-signed certificates with Firefox, an error was seen from Firefox regarding conflicting certificate serial numbers: `sec_error_reused_issuer_and_serial`. |
| **61683** | Fixed an issue with configuring SNMPv3, where an SNMPv3 user with a view could not be added. An error was displayed when attempting to add a name for a view, where the error message stated that view is not a valid reference and the SNMPv3 user with a view could not be added (**Panorama > Collector Groups > Collector Group > Monitoring > Views**). |
| **61682** | An M-100 appliance failed to mount the file system for a RAID Disk Pair. There was no system log generated for this mount failure, and disks were shown as `Available`. This issue has been addressed so that a system log is generated when a disk mount fails, and disk is not made `Available`. |
| **61675** | After creating a Log Forwarding profile using the Panorama web interface (**Device Groups > Objects > Log Forwarding**), the Log Forwarding object always defaulted to being shared across device groups, regardless of whether the **Shared** checkbox had been selected when creating the profile. This has been fixed so that only Log Forwarding objects with **Shared** selected are shared across device groups. Log Forwarding objects where **Shared** is not selected correctly belong only to the Device Group displayed in the Device Group drop-down. |
| **61579** | When session offload is disabled, the sequence and acknowledgment (SEQ/ACK) number check dropped ACK packets. This was due to an issue with the sequence and acknowledgment number check and has been fixed. The workaround for releases prior to PAN-OS 6.0.2 is to use the `set deviceconfig setting tcp asymmetric-path bypass` command. |

| Issue Identifier | Description |
|---|---|
| **61471** | Following a restart, an internal path monitoring failure that showed a false positive triggered a second dataplane restart. This issue has been addressed so that when the dataplane restarts, path monitoring stops and a false positive is not generated so that a second dataplane restart does not occur after reboot. |
| **61437** | On a few occasions, a device stopped authenticating any users. This was due to an issue where bind errors were causing an LDAP file descriptor leak and the issue has been resolved. |
| **61396** | When read using SNMP, the value returned for Dataplane CPU utilization was always zero. This was fixed so that the value returned for dataplane CPU utilization matches the true CPU utilization. |
| **61387** | Data connection of FTP in passive mode sometimes failed on PA-5000 Series devices. This has been fixed so that data connection of FTP in passive mode does not fail. |
| **61366** | When loading a report template, the selected time period was being reset automatically. A fix has been implemented to limit this behavior. |
| **61236** | An issue was resolved to prevent the OSPF MD5 sequence number from resetting during a high availability (HA) failover or routed restart. Before this fix, there was a reset when a failover occurred. This reset caused a longer outage than expected. When upgrading from an earlier PAN-OS 6.0 release version to PAN-OS 6.0.2, allow a five minute grace period before a HA failover or routed restart. Following the five minute grace period, this limitation will not apply to further failovers or upgrades. |
| **61174** | Fixed an issue where license retrieval on the device was failing when the option **Verify Update Server Identity** was enabled on the **Device > Setup > Services** page. |
| **61168** | SYN packets were dropped when a session with the same 5-tuple (same source IP address, destination IP address, source port number, destination port number, and protocol) is received by the firewall at the same time the existing session gets aged-out based on the TIME-WAIT period when source NAT (Dynamic-ip-and-port) was used. |
| **61083** | When using URL Filtering with `ssl-decrypt url-proxy` enabled, the block page failed to display on some browsers using SNI (Server Name Indication) because the firewall failed to retrieve the hostname from the SSL session in order to run decryption. With this fix, the process was improved to retrieve this information correctly. |
| **60985** | Fixed an issue with PA-3000 Series devices where traffic could stop passing through the firewall or the dataplane could restart due to an internal path monitoring failure. |
| **60974** | When a colon was used in password field to generate an API key, using that API key displayed the error `Invalid Credentials`. The way that passwords are handled has been updated to allow a colon as part of the API key. |
| **60928** | In certain situations on the Global Protect gateway, the context for an IPSec tunnel to a GlobalProtect client showed the presence of NAT-T, although there is no NAT between the client and the gateway. Traffic through that IPSec tunnel failed. This issue was caused when a tunnel context ID previously used for a tunnel with NAT-T was re-used for a tunnel without NAT-T. With this fix, tunnel contexts are properly cleared before re-use. |
| **60902** | An issue was resolved where subinterface names with length greater than 16 characters were being truncated in the output of some CLI commands, corresponding web interface sections, and system logs. |


© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **60743** | An untagged subinterface caused stale hardware sessions on a PA-5000 Series device, and traffic failed when a new session was hitting the stale 5 tuple session. This issue has been fixed so that hardware session info is deleted when sessions are closed. |
| **60673** | LDAP groups containing non-ASCII characters in their names could not be processed by the firewall. Policies filtering on these groups were not working properly. With this fix, groups are received by the firewall, regardless of the types of characters used in their names. |
| **60617** | TCP connectivity issues occurred on a virtual machine when SYN Flood protection was enabled with SYN Cookies with the **Activate Rate** set to 0. This was related to an issue where active and passive FTP were not working under different NAT configurations and has been addressed so that both active and passive FTP work correctly under different NAT configurations and do not cause TCP connectivity issues. A workaround for this issue is to set the **Activate Rate** parameter for SYN Flood settings to a value greater than 0 (**Objects > DoS Protection > Flood Protection > SYN Flood**). |
| **60299** | Multicast ICMPv6 traffic was causing a restart after enabling and disabling zone protection profiles multiple times. |
| **60250** | Logging into the Panorama web interface only displayed a blank gray page instead of the Panorama **Dashboard**. This issue has been addressed so that logging into the Panorama web interface correctly displays the Panorama web interface **Dashboard** and management functions. |
| **60036** | Addressed an issue in summary PDF reports where if the **srcuser** field is not in the format of `domain\username`, the **srcuser** field was left empty. |
| **59542** | When sessions were initiated on different dataplanes in a multi-vsys environment, session traffic failed to span across virtual systems. This issue has been resolved so that inter-vsys sessions succeed with a DNAT policy setup. |
| **59289** | An issue has been addressed where dataplane restarts were seen due to PAN-DB becoming unresponsive. New URLs could not be learned and the repeated attempts built up the URL cache and caused an excessive consumption of system resources, making the dataplane unstable. |
| **57736** | Active FTP through NAT (dynamic ip-port) did not work on a virtual machine. This issue has been addressed so that both active and passive FTP work under different NAT configurations. |
| **57600** | Resolved a display issue on the **Monitor > Logs > Configuration** page, where the IPSec pre-shared key was displayed in clear text format in the **After Change** column. The IPSec pre-shared key is no longer displayed in clear text in the Configuration Logs. |
| **57599** | Addressed an issue where the firewall was dropping ARP unicast packets which resulted in incomplete ARP entries and Layer 3 forwarding issues. |
| **57412** | Addressed an issue where User Activity Reports were sometimes not returning any results due to malformed XML output. |

© Palo Alto Networks, Inc.

| Issue Identifier | Description |
|---|---|
| **50798** | Device fans running at lower speeds were sometimes causing a thermal alarm. This occurred when the dataplane restarted while the device was in a temperature alarm state. In this case, the fan's speed control value was set to the default value, causing the fan to run at the incorrect speed. This issue has been addressed so that the fan speed is recalculated based on the device's current temperature following a dataplane restart. |
| **50091** | Management plane services were not performing optimally during peak traffic periods. The size of the ARP table on the management plane was increased to address this issue. |
| **48758** | In some cases, large amounts of compressed HTTP traffic could cause the zip buffer to fill and it would not continue to process additional compressed HTTP traffic. Offload of compressed HTTP traffic now works correctly so that even large amounts of compressed HTTP traffic are processed by the zip buffer. |
| **40883** | Addressed a root partition that was at 100% usage due to an issue where FireMon processes (where FireMon was being used as an external third party monitoring utility to the Palo Alto Networks firewall) continued to run on a device and did not close correctly, even after the FireMon connection had timed out. |
| **24804** | When specifying a default route of 0.0.0.0/0 in an OSPF export rule, the route was advertised with a metric of 255. This was because there was no option to specify a metric when specifying a route in an OSPF export rule. This change removes the metric field from the redistribution profiles, and adds it to the OSPF export rule configuration. |


© Palo Alto Networks, Inc.

# PAN-OS 6.0.1 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.1 release. For new features introduced in PAN-OS 6.0, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Issue Description |
|---|---|
| **61193** | A certificate chain error was causing commit failures due to an IKE code change in 6.0.0. The fix changes the error that causes the commit to fail to a warning so that the commit can be successful. |
| **61004** | Addressed an issue where a user in a custom admin role could only view HIP match logs but was unable to view log details. Clicking the spyglass icon on the web interface opened the **Detailed Log View**, but no information was displayed. |
| **60971** | Following an upgrade to PAN-OS 6.0.0, scheduled Dynamic Updates failed to be pushed from Panorama to managed devices. |
| **60826** | Issues relating to Tags with brackets were occurring after upgrading to PAN-OS 6.0.0, and Tags with brackets could not be created, edited or deleted. In PAN-OS 6.0.1, the brackets in existing Tags are replaced with quotes and new Tags that you create cannot include brackets. |
| **60816** | Following an upgrade to PAN-OS 6.0.0, syslog connection status warnings for all defined syslog connections appeared in the system log every hour and were categorized as critical. This was caused by a scheduled hourly rotation of the syslog-ng log file, during which the syslog-ng daemon would restart. This issue has been fixed by adding a condition to the log file rotation process requiring the log file to be 10 MB or more and the connection status warning will only be seen once every few months. |
| **60780** | On PA-5000 Series firewalls, restarts were seen under heavy load conditions and were due to an internal path monitoring failure. Enhancements have been made to avoid such restarts. |
| **60700** | 1GB copper SFP interfaces belonging to an Aggregate Group would sometimes show the link status as down after a reboot of the firewall. |
| **60677** | Group mapping queries failed when nested groups with long group names were searched. The search queries showed group members to be missing. This was due to the query string (groups/nested groups) being truncated at 39 characters. |
| **60650** | When using the web interface to configure SNMPv3 on a high availability (HA) pair, the **EngineID** field should be optional; however, the **OK** button could be clicked when this field was blank. This has been fixed so that the **EngineID** field is optional. |
| **60564** | On Panorama, a system log message was generated when a log file used for internal processing was truncated. This system log message is not relevant for the administrator and will not be displayed in the system log. |
| **60510** | An attempt to revoke SSL certificates generated on Panorama failed. The issue was caused by an internal verification check failure involving shared certificate objects used in device groups that included firewalls enabled for multiple virtual systems. The verification check has been corrected. |

| Issue Identifier | Issue Description |
|---|---|
| **60505** | SYN packets were dropped if a session with the same source and destination IP addresses and port and IP protocols comes to existing session at the timing of aged-out TIME-WAIT period. |
| **60502** | On the Panorama web interface **Monitor**, the maximum number of pages displayed for logs is not consistent depending on the log numbers displayed per page. |
| **60412** | The user was unable to modify custom logo due to an unrecognized `mime-type` option (`Device > Setup > Operations > Custom Logos`). |
| **60347** | Some service route settings could not be configured when the web interface was set to a language other than English. |
| **60337** | Fixed an issue where XML API showed empty hardware counters for `show interface dedicated-ha1`. |
| **60274** | Fixed an issue where the XML API showed an error when running the `show -availability interface ha1` command using the rest API. |
| **60255** | On occasion, the web interface for the Citrix SDX server fails to execute a reboot or shutdown of the VM-Series firewall. To trigger a reboot or shutdown, use the Force Shutdown or Force Reboot options on the SDX web interface. If you need to modify the instance of the VM-Series firewall, make the changes while the instance is shutdown, and then power it on. |
| **60225** | Fixed an issue where the XML API showed an error when running show session rematch using the rest API. |
| **60201** | In a high availability (HA) active/active configuration, an IPSec key renegotiation timing issue caused the new IPSec session to be set to DISCARD until the next key renegotiation. This caused traffic loss until the next tunnel key renegotiation. |
| **60189** | When high availability (HA) active/passive peers lost communication on HA1 and HA2 links, a race condition caused the dataplane to restart. |
| **60070** | In order for Android devices to be able to connect to the Mobile Security Manager's device check-in for enrollment and subsequent check-ins, the Root CA certificate had to be used to sign the server certificate for the device check-in interface in the portal configuration that was being delivered to Android devices. This applied even if you purchased a server certificate from a well-known, trusted CA as recommended. This was because the GlobalProtect Android app does not look in its system store for certificate verification With this fix, GlobalProtect Android app will first look in its system store for certificate verification. If this fails, the app will proceed to verify certificates against the CAs from the portal configuration. |
| **60063** | SSL decryption with Internet Explorer was inconsistent with certificate common names, resulting in name mismatch errors for some sites. |
| **60035** | When an external zone was configured with a Zone Protection profile applied to it, large IPv6 packets were causing dataplane processes to fail when sent through the zone. |
| **60012** | Addressed issues that were causing filtering on the web interface to return incorrect results. |


© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **60011** | When a User ID Agent Setup template was pushed from Panorama to a managed device, the application content updates were not available for viewing or cloning in the syslog filters list in the web interface (**Device > User Identification > User Mapping > User ID Agent Setup > Syslog Filters**). |
| **59991** | In a shared gateway setup, when the firewall was configured to drop inbound Destination NATed traffic for a particular application, the traffic logs for Destination NATed traffic to be dropped showed the post-translated address instead of the pre-translated address. This issue was fixed by displaying the original source IP address in the traffic logs instead of the translated IP address. |
| **59989** | The Panorama web interface was not displaying data on the **Monitor > Logs > WildFire Submissions** page. A query for the data with no filters continued loading for a significant amount of time and then eventually timed out. This was due to threat logs that were not handled correctly by Panorama running PAN-OS 6.0.0 when received from log collectors running releases previous to PAN-OS 6.0.0. |
| **59973** | For Android devices that were being managed by the GlobalProtect Mobile Security Manager, when authenticating to the GlobalProtect gateway using client certificate authentication, the GlobalProtect app for Android did not look up the identity certificate issued to the device during enrollment. Identity certificates were therefore not being used for gateway authentication as expected. This has been fixed so that Android devices are able to use identity certificates for gateway authentication. |
| **59967** | When the firewall was configured as a GlobalProtect satellite and was receiving access routes from another firewall configured as the GlobalProtect gateway, the routing resource counter for static routes was not incrementing or decrementing correctly. This behavior caused the maximum number of routes to be artificially reached and the firewall stopped accepting routing updates. The fix for this issue readjusts the counters for static routes and total routes when creating a redundant static route or deleting a non-existent one. |
| **59915** | The filter field on the **Policies > Security** page of the web interface returned different results when strings with the same value, but different patterns were entered. For example, xx.xx.xx.0-24 is an address object name which includes the IP/netmask pattern xx.xx.xx.0/24. Entering either string in the filter should return the same search results but one query was showing fewer results than the other. |
| **59890** | When a PA-4050 firewall reached the limit for max supported concurrent decrypted sessions, the dataplane restarted. A fix was added to ensure that device will stop decrypting sessions once the limit is reached and a restart will not occur. |
| **59873** | A TCP session could not be established when SYN Cookies was enabled and when both Aggregate and Classified DoS Protection Profiles were configured. |
| **59772** | Traffic logs from log collectors were not visible on the Panorama web interface. |
| **59707** | NTP information on the firewall was displayed in way that could lead to confusion; for example, stating that the server the device is synced with is not connected (connected: false). NTP information is now displayed more clearly (**Device > Setup > Services**). |
| **59574** | Fixed an issue where an Antivirus profile on Internet Explorer and Firefox browsers was not showing the default action in parenthesis (alert/drop/) for the decoders. |

© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **59471** | Resolved an issue with registering dynamic tags on the VM-Series firewall. The VM-Series firewall did not allow you to register dynamic tags that included the single quote (') character. |
| **59343** | When a security policy was configured that did not have a URL Filtering profile applied to it, URL Filtering logs were still being generated and were visible on the **Monitor > URL Filtering** page. |
| **59309** | User Activity Reports were showing inconsistent results. This was due to the User Activity Report generation taking too long and timing out. The timeout for User Activity Report generation has been extended so that requested reports will run until all data is completely gathered. |
| **59276** | Fixed an issue where botnet reports failed to use URL Filtering data for report generation. |
| **59256** | Performing an SCP import of the logdb file failed with the error: `failed to verify for logdb import`. |
| **59180** | Session setup between a client and server was not completed in a high availability (HA) active/active configuration when configured with multi-vsys and multi-VR. The client from one virtual system could not reach the server located on a second virtual system because the session in the second virtual system was not set up correctly. |
| **59126** | In a high availability (HA) active/passive configuration, OSPF and BGP neighbors went down on an active device after the passive device unexpectedly restarted. |
| **59031** | When admin users tried to login to the CLI without previously logging into the web interface and there was a RADIUS authentication profile configured, the firewall sent out a request to the RADIUS server with an invalid password that was different from the one submitted by the user. This resulted in valid users being unable to authenticate to the RADIUS server. |
| **59018** | A network outage occurred when the number of active sessions reached 100,000 sessions. This occurred on a PA-3000 Series firewall and Detector Threats have been increased on PA-3000 Series firewalls in order to address this issue. |
| **58971** | The CLI output for the `show routing protocol bgp loc-rib-detail` command displayed the community field incorrectly when certain prefix combinations appeared in the IP addresses of the BGP neighbors. |
| **58736** | WildFire email notifications did not contain a date header. |
| **58586** | Fixed the inability to log into the Panorama web interface. This issue occurred because root partition on an M-100 appliance was at 100% usage due to old reports being incorrectly stored in the /tmp directory. |
| **58268** | Traffic was sometimes blocked during SSL decryption when the option **Use OCSP to check certificate status** was enabled. |
| **58215** | The output of the `show routing protocol ospf area` CLI command was updated to provide greater clarity for the values defined. |
| **58212** | New virtual systems could not be added after applying a virtual systems license and could only be added after restarting the management server. |


© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **57997** | In a high availability (HA) active/active configuration, the User-ID agent status was displayed as `connected` on the active-secondary device. This issue has been addressed and the active-secondary device will only show a User-ID agent status of `disconnected` as only the active-primary device connects to the User-ID servers. |
| **57660** | The management ports on a PA-2000 Series firewall did not link-up when connected directly using a straight or crossover cable. |
| **57601** | Fixed an issue where Data Filtering logs were showing incorrect file names when the data pattern was matched against the files. |
| **57261** | A denied session was logged with the Action displayed as `Allow`. This occurred when the application denied was on port 80 and triggered the Captive Portal redirect. |
| **56905** | When a PA-5000 Series firewall received more than 3000 BGP prefixes, the web interface showed an error (`op command for client routed timed out`) when displaying the Local RIB for BGP. Also, when executing the `show routing protocol bgp loc-rib-detail` command, the CLI returned an error (`Server error : op command for client routed timed out`). |
| **55318** | The following commit warning message is now included when nested wildcards are used in a URL Filtering configuration: `Warning: Nested wildcards (*) in URLs may severely impact performance. It is recommended to use a single wildcard to cover multiple tokens or a caret (^) to target a single token.` |
| **50932** | In a high availability (HA) active/passive configuration, `seed` was removed from the `PAN-DB sync seed with HA` message because `MP cache` is what is synchronized instead of `seed`. Also fixed an issue where the synchronized `MP cache` is not loaded on the passive device. |
| **47642** | Logs were not being written to disk because the configuration on the Managed Collector and Collector Group was set up before the Managed Collector established a connection to Panorama. This has been fixed so that Panorama allows you to configure the Collector Group only after the Managed Collector has connected at least once. This allows Panorama to verify the availability of the disk(s) and its size, ensuring that logs are written properly to disk. |
| **45529** | In some User-ID implementations, server session reads picked up capitalized special characters, such as Ü. All capital letters should be set to lower case but this operation was not supported for special characters, which caused a mismatch between group mapping and user mapping. |
| **41347** | Packet capture (PCAP) filters were not filtering information accurately. This fix ensures that the PCAP filters match the criteria defined on the device and accurately capture all relevant frames in the session. |

© Palo Alto Networks, Inc.

© Palo Alto Networks, Inc.

# PAN-OS 6.0.0 Addressed Issues

The following table lists the issues that are fixed in the PAN-OS 6.0.0 release. For new features, associated software versions, known issues, and changes in default behavior, see PAN-OS 6.0 Release Information.

| Issue Identifier | Issue Description |
|---|---|
| **60347** | Some service route settings could not be configured when the web interface was set to a language other than English. |
| **59772** | Traffic logs from log collectors are not visible on the Panorama web interface. |
| **59707** | NTP information on the firewall was displayed in way that could lead to confusion; for example, stating that the server the device is synced with is not connected (connected: false). NTP information is now displayed more clearly. |
| **59407** | NetFlow (type 4) messages were appearing in the traffic log database and reports. |
| **59128** | After logging in to Panorama using the CLI with RADIUS credentials, the following error message was printed: `Server error : show -> system -> setting -> multi-vsys is unexpected`. |
| **59031** | When admin users tried to login to the CLI without previously logging into the web interface and there was a RADIUS authentication profile configured, the firewall sent out a request to the RADIUS server with an invalid password that was different from the one submitted by the user. This resulted in valid users being unable to authenticate to the RADIUS server. |
| **59030** | Certificates generated during SSL decryption were not adhering to the ASN.1 format. This was leading to the SSL connection being dropped by some servers. |
| **58885** | The `test nat-policy-match` command now properly displays results for no-nat rules. |
| **58736** | WildFire email notifications did not contain a date header. |
| **58733** | The fields in the CSV report were displayed incorrectly after performing a CSV export on the **Monitor > HIP Match** page on the Panorama web interface. |
| **58614** | Local users discovered by WMI query were mapped as the local user of the computer, instead of Unknown as is the expected behavior. |
| **58347** | Suppressed extraneous messages (for example, disabling of an interrupt request that occurs within the underlying subsystem) from displaying on the console. These messages are now logged in the system log only. |
| **58264** | Previously, the `debug software virt-limit limit` command showed an incorrect max value: `4294967295`. The max value has been fixed to display in kilobytes. |
| **58223** | Captive portal was not presenting a complete certificate chain to the client. It presented only the end certificate and not the intermediate certificate. |
| **58215** | The output from the `show routing protocol ospf area` CLI command was rearranged to provide greater clarity in the values defined. |

| Issue Identifier | Issue Description |
|---|---|
| **57975** | It was not possible using Panorama to proxy a REST API call for retrieving report information from a firewall. |
| **57960** | When the Palo Alto Networks firewall was configured to support several virtual systems, the firewall administrator could not revert the **Destination Interface** in a NAT Policy Rule back to the option any after an interface had been selected. This was because the any field in the NAT to-interface configuration had an incorrect schema value. The incorrect schema was fixed by adding any as a default NAT to-interface value in the configuration. |
| **57927** | When authenticating through captive portal, there was a delay after the authentication redirect for Firefox and Chrome browsers. This has been corrected by closing the socket after the redirect. |
| **57874** | DNS resolution did not turn off when the **Resolve Hostname** checkbox was cleared in the **Monitor** tab, and the Palo Alto firewall continued to display the hostnames instead of the IP addresses. IP addresses are now displayed when the **Resolve Hostname** checkbox is cleared. |
| **57768** | A DHCP server did not differentiate between DHCP Clients when the DHCP Client Identifier in the DHCP request exceeded 32 bytes. The maximum size of the DHCP Client Identifier has been increased to 312 bytes. |
| **57660** | PA-2000 Series platform management ports did not link-up when connected directly using a straight or cross cable. |
| **57608** | When using multiple NetFlow hosts across multiple profiles, instances of the FlowSequence number were skipped. The expected behavior is that the value is cumulative, and should be used by the Collector to identify whether any Export Packets have been missed. |
| **57535** | Fixed an issue where the user was not able to create a QoS profile with an egress bandwidth greater than 50 Mbps on a virtual firewall (**Network > Network Profiles > QoS Profile**). |
| **57507** | The option **L3 Forwarding Enabled** in the configuration of a VLAN has been removed. In pre-6.0 releases, enabling or disabling this option did not have an effect on traffic forwarding. Enabling or disabling L3 forwarding on a VLAN should be performed by adding or removing an L3 VLAN interface to the VLAN configuration. |
| **57448** | The **IRC** checkbox in the Botnet Configuration window (**Monitor > Botnet**) was not displayed on the web interface when the language was set to Japanese and a Chrome browser was being used. |
| **57360** | CLI help for `show session all filter destination` command is showing `<ip/netmask>` instead of `<ip>`. |
| **57258** | Both HTTP and HTTPS were available when accessed directly from the management interface; however, HTTP was unavailable when accessed using a subinterface. |
| **57159** | The dataplane was passing traffic even though the management plane was rebooted and could not boot. |
| **57154** | On a PA-5000 series firewall, the QoS rate is adjusted slightly to accommodate hardware limitations. The following help message now is displayed on the configuration window on the web interface: `Bandwidth limits shown include hardware adjustment factor`. |


© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **57098** | In some Layer 2 configurations, multicast traffic passing through the firewall was resulting in both forward and drop counters incrementing due to the packets being broadcast. Additionally, the multicast packet was included in both the transmit and drop stage dataplane packet capture. New global counters were added to clarify the actions being taken by the firewall when processing multicast packets in a Layer 2 configuration. |
| **56905** | When a PA-5000 Series firewall received more than 3000 BGP prefixes, the web interface showed an error when displaying the Local RIB for BGP: `op command for client routed timed out`. Additionally, when the `show routing protocol bgp loc-rib-detail` command was issued, the CLI returned the error: `Server error : op command for client routed timed out`. |
| **56858** | A cache corruption prevented the user from downloading files when clicking the Continue button in the File Blocking Continue page. |
| **56802** | In a single-vsys setup, a Log Forwarding Profile created on the web interface was not displayed after issuing the `show shared log-settings profile` CLI command. |
| **56787** | After an upgrade, the captive portal custom response page shows `::ffff:` before the IP address. |
| **56703** | In the web interface, global timeout values were displayed in addition to the application level timeout values that actually took effect. This has been updated to show only application level timeout values. |
| **56367** | Fixed an issue where NetFlow data could not be exported for all subinterface types. NetFlow records were not picked up by the log-receiver. |
| **56107** | Addressed dataplane restarts that occurred intermittently on the PA-3000 Series devices deployed in a high availability (HA) configuration. |
| **56087** | Log collectors were optimized in PAN-OS 6.0 for quicker failover and failback. |
| **55833** | GRE port information was not mapped correctly on the VM-Series platforms, causing predict sessions to not match and leading to dropped packets. |
| **55774** | On the web interface, setting the value for max-rows-in-csv-export did not work when it was set to more than 65535. |
| **55696** | Misspellings were displayed in the output for the `set session processing-cpu` command. The misspellings have been corrected. |
| **55693** | Added an enhancement to reduce the routed log in order to help reduce OSPF flaps. |
| **55407** | User-ID virtual memory was exceeding its limit in a multi-vsys environment when a large number of LDAP objects were returned to the firewall. With this fix, LDAP queries made by the firewall will filter on groups specified in the include-list. |
| **55387** | When using local user groups to assign users to particular gateways, the connection to the Global Protect server for the users in that local group failed. |
| **55342** | When an LDAP Server Profile was configured with multiple LDAP servers and/or multiple User-ID agents were configured to be used as LDAP proxy servers and the primary server was not connected, the firewall continually attempted to connect to the secondary LDAP server (or LDAP proxy server) but was unable to establish a connection. |

© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **55111** | Fixed an issue where a FIN packet was dropped when the sequence number was out of sync after traffic triggered session reuse and was offloaded. With this fix, the sequence-number check for offloaded reused sessions is skipped (because the data plane processor cannot track sequence numbers after traffic is offloaded). |
| **54958** | Upon opening a PCAP on the firewall, escape sequences were displayed instead of the special characters in data part. A fix is provided to display the characters correctly. |
| **54949** | A commit failed when DHCPv6 relay was configured on an interface that did not have an IPv4 address. |
| **54755** | An issue was addressed where creating a static route with the next hop set to None and cloning it or going back into it was changing the next hop settings to <IP Address> from None. |
| **54676** | In the web interface, on the **Device > User Identification > Group Mapping Settings > Group Mapping > Group Include List** tab, the list of Available Groups to add to the Included Group list displayed approximately the first 200 groups, with the option to select **more** to view more group entries. However, clicking **more** failed to display more group entries, even when several more groups are defined and should be available. |
| **54547** | Fixed an issue where high availability (HA) peer HA2 IP information was not getting updated after issuing the `show high-availability all` CLI command. |
| **54486** | Added support for both single quote and double quote values when entering options using the Command Line Interface (CLI). |
| **54283** | An auto commit failed during a threat database update, displaying the error `Threat database handler failed`. |
| **54265** | The system log message `Antivirus job failed` has been updated and the following will be reported in the system log instead: `Antivirus update job failed`. |
| **54113** | A Forwarding Information Base (FIB) table entry discrepancy caused SSH packets to be sent back. This only occurred on PA-2000 Series firewalls. |
| **53888** | On PA-5000 Series devices, the DIPP limit was causing the following system error when trying to add more NAT policies to the firewall: `Error: Number of dynamic-ip-and-port rules (251) exceeds vsys capacity (250) Error: Failed to parse nat policy The maximum number of DIPP has now been increased`. |
| **53632** | Fixed a BGP aggregate policy issue where the aggregate route was no longer advertised when a more specific prefix within the aggregate range was learned. |
| **53615** | When enabling IPv6 on an interface, link local IPv6 routes were counted towards the rtm_total/connected/ipv6; however, the Link Local IPv6 routes were not installed to the Forward Information Base (FIB) on the dataplane. |
| **53554** | Disks in a Panorama VM OVF were misaligned with NetApp and caused performance degradation with some storage devices. |
| **53514** | A high availability (HA) active/active configuration for IPv6 using Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) did not behave consistently when stateless address autoconfiguration (SLAAC) was also configured. |


© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **53148** | Output of `debug dataplane packet-diag show setting` command truncates the interface name to 15 characters. |
| **53059** | Role-based admin users without privileges to access logs or the **Monitor** tab were able to view logs using the **Dashboard** widgets. |
| **52847** | Link monitoring and Path monitoring were on hold when a commit started and until one minute after the commit was done. Changes are introduced to remove the hold on the Link and Path monitoring during Phase -1 of commits. |
| **52777** | Link and Path monitoring were not always working properly during the commit process. |
| **52738** | Reset was sent to Captive Portal clients when trying to load multiple pages before logging in to the portal. |
| **52629** | PAN-DB reverted back to BrightCloud due lack of management connection for first reboot. |
| **52567** | The loading icon was not shown when using the list of users to add a source user to a security policy on the web interface. |
| **52214** | Some traffic was getting dropped if the number of routes in the routing table was high. |
| **52184** | Changing the Jumbo Frame settings on the device without restarting the entire device caused the dataplane to experience an unexpected restart. This has since been fixed so that when you change Jumbo Frame settings, an entire device reboot is no longer required and a dataplane restart will work. |
| **52128** | Fixed an issue where a management profile was configured on an interface and the clients were not getting IP addresses from the DHCP server when the device was configured as a DHCP relay agent. |
| **52050** | After manually upgrading PAN-OS, no **Reboot** button was visible, as it was in previous releases. A message was displayed instead that the user must reboot the device by closing the current window and then rebooting. |
| **51955** | The CLI displayed two counters listed under `IPv6 filter`, even though they also applied to IPv4. A change was made to list them under `IPv(4/6) filter`. |
| **51880** | Dynamic role-based device administrators did not have the ability to save, export, load, and revert a configuration on the firewall or Panorama. This fix provides these capabilities to the administrators. |
| **51824** | Device Groups added to multiple virtual systems were not always shown as managed devices on the web interface (**Panorama > Device Groups > Device Group**). |
| **51648** | In a high availability (HA) active/passive configuration, if NAT exists for outbound FTP connections and the interface IP address is used for the NAT, the ftp-data session would not synchronize to the passive device. |
| **51597** | When the XML API was used to push IP address, port range and username information to a firewall deployed in high availability (HA) configuration, the details were not synchronized with the HA peer. |

© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **51091** | Two-factor authentication (where both a client certificate profile and an authentication profile are configured) was not functioning as expected. The client was not required to provide the login credentials associated with the authentication profile after successfully authenticating with the client certificate. |
| **51089** | Fixed an issue where repeat count in threat logs was resulting in incorrect values. |
| **51062** | Inter-vsys sessions that traverse the firewall and terminate on a firewall interface would fail. This has been fixed. |
| **51042** | Certificates that were generated prior to master key changes could continue to be used. |
| **51000** | On a redundant Power Supply system on a PA-5000 Series device, there was no system log visible when removing or adding redundant Power Supply. Logging for these events has been added. |
| **50963** | Panorama software deployment failed to deploy when the **OK** button was clicked. |
| **50936** | Crypto cores were created when a SIGTERM signal was received while the management plane was starting. |
| **50817** | Fixed an issue for firewalls running PAN-OS 4.1.6 and later releases where GlobalProtect users were unable to connect to the GlobalProtect gateway when the external-facing interface on the gateway was configured with dynamic PPPoE and a loopback interface was configured for the destination interface to the GlobalProtect portal. With this fix, you can configure a loopback interface for GlobalProtect even when the external-facing interface for the gateway is dynamic. |
| **50606** | Captive Portal authentication failed when the username contained the character &. This issue has been addressed so that & is a valid character and Captive Portal authentication is successful when a username contains the character &. |
| **50478** | The certificate signing request (CSR) generated by the firewall had a Challenge Attribute set by default. If configured, the signing entity could use this attribute or ignore it. Since this attribute was not being ignored by some signing entities, the behavior has been updated so that the Challenge Attribute is not set by default. |
| **50310** | A destination based service route for DNS prevented an FQDN query from refreshing. |
| **50091** | A possible memory leak caused management plane services to not perform optimally during peak traffic periods. |
| **50079** | Added logging enhancements in order to help identify root cause. |
| **50048** | The `show session all filter from <zone> to <zone>` CLI command displayed no active sessions when there were active sessions that should have been displayed in the output. |
| **49851** | In PAN-OS 6.0, DoS enforcement is now performed in CPU prior to session installation. |
| **49828** | In custom reports, source and destination country are now available in the Query Builder as grouping options to organize the report. |
| **49727** | Navigating to the **Network > Interface > Ethernet** tab took 12-15 seconds for the screen to populate the interface data. |


© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **49294** | The **ACC** (Application Command Center) tab on the Panorama web interface failed to display complete sections and appeared to be stalled, showing the error message: `3 requests sent 1 response received`. |
| **49038** | Time zones were not automatically converted for Dynamic Update package release times. |
| **49015** | Fixed a dataplane restart issue that occurred when Jumbo Frames were enabled and the packets received buffer was high. |
| **48896** | In rare cases, abrupt restarting (for example, a power outage) lead to internal system file corruption. This was related to checking OS image integrity and cannot upgrade. Preventative measures were put in place to prevent issues before and after the internal file updating. |
| **48729** | In Panorama, disabling the **Share Unused Address and Service Objects with Devices** feature returned an error stating that the shared address is not a valid reference. This occurred when a non-shared address group—that was assigned to a specific device group—contained a shared address or an address group was pushed. This issue has been fixed so that such a configuration is supported. |
| **48709** | Fixed an issue where setting a packet capture filter in the web interface did not work until the filter was reset by removing the automatically added 0.0.0.0. |
| **48703** | This fixes a NAT pool leak issue when a SYN packet on TCP/443 was sent to an address on an interface on which GlobalProtect was configured but which was not its primary address. A NAT port was allocated, the connection failed, and the session was freed, but the allocated NAT port is not cleared. |
| **48584** | On Panorama, there were long delays committing a policy due to the option **Share Unused Address and Service Objects with Devices** being cleared in large configurations. The delay was introduced as the system performs a calculation of the unused objects on commit. Commit times have been improved for large configurations. |
| **48093** | Configured address objects were not displayed as resolved on the Panorama web interface. On both the **ACC** tab and the **Monitor > Logs > Traffic** tab, host names defined in the address objects were not displayed, and the IP address was shown in the Host Name columns. |
| **47616** | Devices which were no longer managed devices (had been managed devices previously but were not anymore) were displayed on the **Panorama > Device Deployment > Licenses** page on the Panorama web interface. |
| **47461** | Fixed an issue where SIP sessions were going into offload state after a content installation causing SIP connectivity issues. |
| **47071** | In PAN-OS 6.0.0, you can now rename and push a shared object from Panorama to a managed firewall if you used that shared object in a local policy. |
| **47007** | An enhanced mechanism to hold control session packets being sent out before predict session is now installed on the master dataplane. |
| **46535** | When using an Internet Explorer browser and a Block / Continue page appears when attempting to download a file, clicking the Continue option did not download the file. |

© Palo Alto Networks, Inc.

| Issue Identifier | Issue Description |
|---|---|
| **46308** | The full User-ID Mapping table is now synchronized between peers in a high availability (HA) cluster. |
| **46134** | On the Panorama web interface, DHCP server settings displayed for entries on the **Network > DHCP** page were not displayed on the DHCP Server window that is displayed when clicking on one of the specific DHCP server entries. |
| **45529** | In some User-ID implementations, server session reads picked up capitalized special characters such as Ü. Normally all capitals are set to lower case, but this operation was not supported for special characters, causing a mismatch between group mapping and ip mapping. |
| **44925** | When a Virtual Router interface was deleted, added, or updated with a new IP/mask, all local Virtual Router interfaces on the management plane were uninstalled and then reinstalled. With this fix, the management plan will assess if all Virtual Router interfaces change before automatically uninstalling and reinstalling them all; the management plane will not continue to uninstall and reinstall all Virtual Router interfaces unless they have all been changed. |
| **43280** | Prior to PAN-OS 6.0.0, NetFlow data could not be exported on a per subinterface basis. Starting in PAN-OS 6.0.0, NetFlow data can be exported on a per subinterface basis. |
| **41472** | When a DNS Proxy object was configured with static entries, hostnames assigned to the DNS Proxy were resolved as expected to the IP addresses listed on the Static Entries tab (**Network > DNS Proxy**). However, when setting the DNS Proxy Object as the DNS Service on the **Device > Setup > Services** dialog, all DNS queries from the management interface ignored the defined static entries. |
| **40648** | Validation logic has been added to PAN-OS software image files to prevent upgrade failures due to file corruption. |
| **39368** | Enhancements have been made to the web interface so that high availability (HA) link status is displayed with green or red indicators on the **High Availability** widget on the **Dashboard** tab. A green indicator signifies that the link is up on the HA port and heartbeats (keep-alive messages) are being sent and received. A red indicator signifies that the link on the HA port is down or that heartbeats (keep-alive messages) are not being received at all. (For HA3 interfaces, the green and red indicators signify only whether the link is up or down.) |


© Palo Alto Networks, Inc.

# Getting Help

The following topics provide information on where to find out more about our products and how to request support:


- Related Documentation
- Requesting Support


## Related Documentation

Refer to the following documents on the Technical Documentation portal at https://www.paloaltonetworks.com/documentation for more information on our products:

- New Features Guide—Detailed information on configuring the features introduced in this release.
- PAN-OS Administrator's Guide—Provides the concepts and solutions to get the most out of your Palo Alto Networks next-generation firewalls. This includes taking you through the initial configuration and basic set up on your Palo Alto Networks firewalls.
- Panorama Administrator's Guide—Provides the basic framework to quickly set up the Panorama virtual appliance or the M-100 appliance for centralized administration of the Palo Alto Networks firewalls.
- WildFire Administrator's Guide—Provides information on deploying, operating, and maintaining the WildFire cloud and the WildFire WF-500 appliance and the Palo Alto Networks firewalls.
- VM-Series Deployment Guide—Provides details on deploying and licensing the VM-Series firewall on all supported hypervisors. It includes example of supported topologies on each hypervisor.
- GlobalProtect Administrator's Guide—Takes you through the configuration and maintenance of your GlobalProtect infrastructure.
- Online Help System—Detailed, context-sensitive help system integrated with the firewall web interface.
- Open Source Software (OSS) Listings—OSS licenses used with Palo Alto Networks products and software:
  - PAN-OS 6.0
  - Panorama 6.0
  - WildFire 6.0

© Palo Alto Networks, Inc.

## Requesting Support

For contacting support, for information on support programs, to manage your account or devices, or to open a support case, refer to https://www.paloaltonetworks.com/support/tabs/overview.html.

To provide feedback on the documentation, please write to us at: documentation@paloaltonetworks.com.

## Contact Information

**Corporate Headquarters:**

Palo Alto Networks
4401 Great America Parkway
Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-us

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2014–2016 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at http://www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

**Revision Date: February 23, 2016**

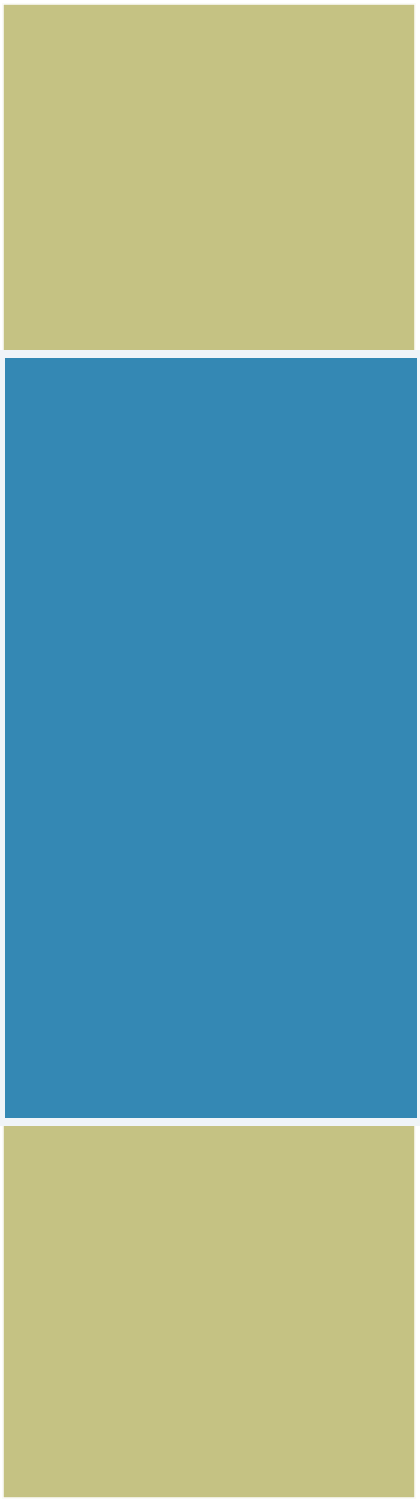

Palo Alto Networks

# PAN-OS® Administrator's Guide
Version 6.1

## Contact Information

**Corporate Headquarters:**
Palo Alto Networks
4401 Great America Parkway
Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-us

## About this Guide

This guide takes you through the configuration and maintenance of your Palo Alto Networks next-generation firewall. For additional information, refer to the following resources:

- For information on the additional capabilities and for instructions on configuring the features on the firewall, refer to https://www.paloaltonetworks.com/documentation.
- For access to the knowledge base, discussion forums, and videos, refer to https://live.paloaltonetworks.com.
- For contacting support, for information on support programs, to manage your account or devices, or to open a support case, refer to https://www.paloaltonetworks.com/support/tabs/overview.html.
- For the most current PAN-OS 6.1 release notes, go to https://www.paloaltonetworks.com/documentation/61/pan-os/pan-os-release-notes.html.

To provide feedback on the documentation, please write to us at: documentation@paloaltonetworks.com.

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2014–2016 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at http://www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

**Revision Date: March 23, 2016**

# Table of Contents


## Getting Started . . . . . 13

Integrate the Firewall into Your Management Network . . . . . 14
- Determine Your Management Strategy . . . . . 14
- Perform Initial Configuration . . . . . 15
- Set Up Network Access for External Services . . . . . 17
- Register the Firewall . . . . . 21
- Activate Licenses and Subscriptions . . . . . 21
- Manage Content Updates . . . . . 23
- Install Software Updates . . . . . 26

Create the Security Perimeter . . . . . 28
- Basic Interface Deployments . . . . . 28
- About Security Policy . . . . . 30
- Plan the Deployment . . . . . 31
- Configure Interfaces and Zones . . . . . 32
- Set Up Basic Security Policies . . . . . 34

Enable Basic Threat Prevention Features . . . . . 41
- Enable WildFire . . . . . 41
- Scan Traffic for Threats . . . . . 43
- Control Access to Web Content . . . . . 47

Best Practices for Completing the Firewall Deployment . . . . . 51

## Device Management . . . . . 53

Management Interfaces . . . . . 54
- Use the Web Interface . . . . . 55
- Use the Command Line Interface (CLI) . . . . . 60
- Use the XML API . . . . . 62

Manage Firewall Administrators . . . . . 64
- Administrative Roles . . . . . 64
- Administrative Authentication . . . . . 65
- Create an Administrative Account . . . . . 65

Reference: Web Interface Administrator Access . . . . . 69
- Web Interface Access Privileges . . . . . 69
- Panorama Web Interface Access . . . . . 105

Reference: Port Numbers Used by Palo Alto Networks Devices . . . . . 108
- Ports Used for Management Functions . . . . . 108
- Ports Used for HA . . . . . 109
- Ports Used for Panorama . . . . . 110
- Ports Used for User-ID . . . . . 111

Reset the Firewall to Factory Default Settings . . . . . 112


© Palo Alto Networks, Inc.

## Certificate Management . . . . . 113

Keys and Certificates . . . . . 114

Certificate Revocation . . . . . 116
- Certificate Revocation List (CRL) . . . . . 116
- Online Certificate Status Protocol (OCSP) . . . . . 117

Certificate Deployment . . . . . 118

Set Up Verification for Certificate Revocation Status . . . . . 119
- Configure an OCSP Responder . . . . . 119
- Configure Revocation Status Verification of Certificates Used for User/Device Authentication . . . 120
- Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption . . . . . 121

Configure the Master Key . . . . . 123

Obtain Certificates . . . . . 124
- Create a Self-Signed Root CA Certificate . . . . . 125
- Generate a Certificate on the Firewall . . . . . 126
- Import a Certificate and Private Key . . . . . 127
- Obtain a Certificate from an External CA . . . . . 128

Configure a Certificate Profile . . . . . 130

Configure the Key Size for SSL Forward Proxy Server Certificates . . . . . 133

Revoke and Renew Certificates . . . . . 134
- Revoke a Certificate . . . . . 134
- Renew a Certificate . . . . . 134

Secure Keys with a Hardware Security Module . . . . . 135
- Set up Connectivity with an HSM . . . . . 136
- Encrypt a Master Key Using an HSM . . . . . 142
- Store Private Keys on an HSM . . . . . 144
- Manage the HSM Deployment . . . . . 146

## High Availability . . . . . 147

HA Overview . . . . . 148

HA Concepts . . . . . 149
- HA Modes . . . . . 149
- HA Links and Backup Links . . . . . 149
- Device Priority and Preemption . . . . . 151
- Failover Triggers . . . . . 152
- HA Timers . . . . . 152

Set Up Active/Passive HA . . . . . 155
- Prerequisites for Active/Passive HA . . . . . 156
- Configuration Guidelines for Active/Passive HA . . . . . 157
- Configure Active/Passive HA . . . . . 159
- Define HA Failover Conditions . . . . . 165
- Verify Failover . . . . . 166

HA Resources . . . . . 167



© Palo Alto Networks, Inc.

## Reports and Logging . . . . . 169

Use the Dashboard . . . . . 170

Use the Application Command Center . . . . . 172
- ACC Risk Level . . . . . 173
- ACC Charts . . . . . 174
- ACC Detail Pages . . . . . 176
- Use the ACC . . . . . 177

Use App-Scope . . . . . 178
- Summary Report . . . . . 179
- Change Monitor Report . . . . . 180
- Threat Monitor Report . . . . . 182
- Threat Map Report . . . . . 183
- Network Monitor Report . . . . . 184
- Traffic Map Report . . . . . 185

Take Packet Captures . . . . . 186

Monitor the Firewall . . . . . 188
- Monitor Applications and Threats . . . . . 189
- Monitor Log Data . . . . . 190
- Monitor the Dashboard . . . . . 194
- View Reports . . . . . 195

Forward Logs to External Services . . . . . 196
- Define Remote Logging Destinations . . . . . 197
- Enable Log Forwarding . . . . . 204
- Log Forwarding Profiles . . . . . 205

Monitor the Firewall Using SNMP . . . . . 208

Monitor the Firewall Using NetFlow . . . . . 210

NetFlow Templates . . . . . 211

Identify Firewall Interfaces in External Monitoring Systems . . . . . 214

Manage Reporting . . . . . 216
- About Reports . . . . . 217
- View Reports . . . . . 218
- Disable Predefined Reports . . . . . 219
- Generate Custom Reports . . . . . 220
- Generate Botnet Reports . . . . . 226
- Manage PDF Summary Reports . . . . . 228
- Generate User/Group Activity Reports . . . . . 230
- Manage Report Groups . . . . . 231
- Schedule Reports for Email Delivery . . . . . 232

Syslog Field Descriptions . . . . . 233

## User-ID . . . . . 245

User-ID Overview . . . . . 246

User-ID Concepts . . . . . 248
- Group Mapping . . . . . 248
- User Mapping . . . . . 248

Enable User-ID . . . . . 252

Map Users to Groups . . . 253
Map IP Addresses to Users . . . 255
Configure User Mapping Using the Windows User-ID Agent . . . 256
Configure User Mapping Using the PAN-OS Integrated User-ID Agent . . . 262
Configure User-ID to Receive User Mappings from a Syslog Sender . . . 265
Map IP Addresses to User Names Using Captive Portal . . . 275
Configure User Mapping for Terminal Server Users . . . 281
Send User Mappings to User-ID Using the XML API . . . 289
Configure a Firewall to Share User Mapping Data with Other Firewalls . . . 290
Enable User- and Group-Based Policy . . . 292
Verify the User-ID Configuration . . . 295

App-ID . . . 299
App-ID Overview . . . 300
Manage Custom or Unknown Applications . . . 301
Best Practices for Using App-ID in Policy . . . 303
Applications with Implicit Support . . . 304
Application Level Gateways . . . 307
Disable the SIP Application-level Gateway (ALG) . . . 308

Threat Prevention . . . 309
Set Up Security Profiles and Policies . . . 310
Set Up Antivirus, Anti-spyware, and Vulnerability Protection . . . 311
Set Up Data Filtering . . . 313
Set Up File Blocking . . . 317
Prevent Brute Force Attacks . . . 319
Customize the Action and Trigger Conditions for a Brute Force Signature . . . 320
Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions . . . 323
Enable Passive DNS Collection for Improved Threat Intelligence . . . 326
Use DNS Queries to Identify Infected Hosts on the Network . . . 327
DNS Sinkholing . . . 328
Configure DNS Sinkholing . . . 329
Identify Infected Hosts . . . 333
Content Delivery Network Infrastructure for Dynamic Updates . . . 335
Threat Prevention Resources . . . 337

Decryption . . . 339
Decryption Overview . . . 340
Decryption Concepts . . . 341
Keys and Certificates for Decryption Policies . . . 342
SSL Forward Proxy . . . 344
SSL Inbound Inspection . . . 346
SSH Proxy . . . 347
Decryption Exceptions . . . 348
Decryption Port Mirroring . . . 349


© Palo Alto Networks, Inc.

Configure SSL Forward Proxy . . . 350
Configure SSL Inbound Inspection . . . 355
Configure SSH Proxy . . . 357
Configure Decryption Exceptions . . . 358
Exclude Traffic From Decryption . . . 358
Exclude a Server From Decryption . . . 359
Enable Users to Opt Out of SSL Decryption . . . 360
Configure Decryption Port Mirroring . . . 362
Temporarily Disable SSL Decryption . . . 365

URL Filtering . . . 367

URL Filtering Overview . . . 368
URL Filtering Vendors . . . 368
Interaction Between App-ID and URL Categories . . . 369
URL Filtering Concepts . . . 370
URL Categories . . . 371
URL Filtering Profile . . . 372
URL Filtering Profile Actions . . . 373
Block and Allow Lists . . . 374
Safe Search Enforcement . . . 375
Container Pages . . . 377
HTTP Header Logging . . . 378
URL Filtering Response Pages . . . 379
URL Category as Policy Match Criteria . . . 381
PAN-DB Categorization Workflow . . . 382
PAN-DB URL Categorization Components . . . 382
PAN-DB URL Categorization Workflow . . . 384
Enable a URL Filtering Vendor . . . 385
Enable PAN-DB URL Filtering . . . 386
Enable BrightCloud URL Filtering . . . 387
Determine URL Filtering Policy Requirements . . . 389
Monitor Web Activity . . . 391
Interpret the URL Filtering Logs . . . 392
Use the ACC to Monitor Web Activity . . . 394
View URL Filtering Reports . . . 395
View the User Activity Report . . . 396
Configure Custom URL Filtering Reports . . . 398
Configure URL Filtering . . . 399
Log Client IP Addresses for Proxied HTTP/HTTPS Requests . . . 401
Replace Source IP Addresses with XFF Values in URL Filtering Logs . . . 401
Add XFF Values to URL Filtering Logs . . . 401
Customize the URL Filtering Response Pages . . . 403
Configure URL Admin Override . . . 404

© Palo Alto Networks, Inc.

Enable Safe Search Enforcement . . . . . 406
Block Search Results that are not Using Strict Safe Search Settings . . . . . 406
Enable Transparent Safe Search Enforcement . . . . . 410
URL Filtering Use Case Examples . . . . . 414
Use Case: Control Web Access . . . . . 415
Use Case: Use URL Categories in Policy . . . . . 419
Troubleshoot URL Filtering . . . . . 421
Problems Activating PAN-DB . . . . . 421
PAN-DB Cloud Connectivity Issues . . . . . 422
URLs Classified as Not-Resolved . . . . . 423
Incorrect Categorization . . . . . 424
URL Database Out of Date . . . . . 425

Quality of Service . . . . . 427

QoS Overview . . . . . 428
QoS Concepts . . . . . 430
QoS for Applications and Users . . . . . 430
QoS Profile . . . . . 431
QoS Classes . . . . . 432
QoS Policy . . . . . 433
QoS Egress Interface . . . . . 434
QoS Cleartext and Tunneled Traffic . . . . . 435
Configure QoS . . . . . 436
Configure QoS for a Virtual System . . . . . 441
QoS Use Case Examples . . . . . 448
Use Case: QoS for a Single User . . . . . 449
QoS for Voice and Video Applications . . . . . 452

VPNs . . . . . 455

VPN Deployments . . . . . 456
Site-to-Site VPN Overview . . . . . 457
Site-to-Site VPN Concepts . . . . . 458
IKE Gateway . . . . . 458
Tunnel Interface . . . . . 458
Tunnel Monitoring . . . . . 459
Internet Key Exchange (IKE) for VPN . . . . . 459
Set Up Site-to-Site VPN . . . . . 463
Set up an IKE Gateway . . . . . 463
Define Cryptographic Profiles . . . . . 465
Set up an IPSec Tunnel . . . . . 467
Set up Tunnel Monitoring . . . . . 470
Test VPN Connectivity . . . . . 472
Interpret VPN Error Messages . . . . . 472
Site-to-Site VPN Quick Configs . . . . . 474
Site-to-Site VPN with Static Routing . . . . . 474
Site-to-Site VPN with OSPF . . . . . 478
Site-to-Site VPN with Static and Dynamic Routing . . . . . 483


© Palo Alto Networks, Inc.

## Large Scale VPN (LSVPN) . . . . . 491

LSVPN Overview . . . . . 492
Create Interfaces and Zones for the LSVPN . . . . . 493
Enable SSL Between GlobalProtect LSVPN Components . . . . . 495
About Certificate Deployment . . . . . 495
Deploy Server Certificates to the GlobalProtect LSVPN Components . . . . . 495
Configure the Portal to Authenticate Satellites . . . . . 499
Configure GlobalProtect Gateways for LSVPN . . . . . 501
Prerequisite Tasks . . . . . 501
Configure the Gateway . . . . . 501
Configure the GlobalProtect Portal for LSVPN . . . . . 504
Prerequisite Tasks . . . . . 504
Configure the Portal . . . . . 505
Define the Satellite Configurations . . . . . 506
Prepare the Satellite Device to Join the LSVPN . . . . . 509
Verify the LSVPN Configuration . . . . . 512
LSVPN Quick Configs . . . . . 513
Basic LSVPN Configuration with Static Routing . . . . . 514
Advanced LSVPN Configuration with Dynamic Routing . . . . . 517

## Networking . . . . . 521

Interface Deployments . . . . . 522
Virtual Wire Deployments . . . . . 522
Layer 2 Deployments . . . . . 525
Layer 3 Deployments . . . . . 525
Tap Mode Deployments . . . . . 526
Virtual Routers . . . . . 527
Static Routes . . . . . 529
RIP . . . . . 531
OSPF . . . . . 533
OSPF Concepts . . . . . 533
Configure OSPF . . . . . 536
Configure OSPFv3 . . . . . 541
Configure OSPF Graceful Restart . . . . . 544
Confirm OSPF Operation . . . . . 545
BGP . . . . . 548
Session Settings and Timeouts . . . . . 553
Transport Layer Sessions . . . . . 554
TCP . . . . . 554
UDP . . . . . 556
ICMP . . . . . 556
Configure Session Timeouts . . . . . 557
Configure Session Settings . . . . . 559


© Palo Alto Networks, Inc.

DHCP . . . 562
DHCP Overview . . . 562
DHCP Messages . . . 563
DHCP Addressing . . . 563
DHCP Options . . . 565
Restrictions to Palo Alto Networks Implementation of DHCP . . . 566
Configure an Interface as a DHCP Server . . . 566
Configure an Interface as a DHCP Client . . . 570
Configure an Interface as a DHCP Relay Agent . . . 571
Monitor and Troubleshoot DHCP . . . 571
NAT . . . 573
Purpose of NAT . . . 573
NAT Rules and Security Policies . . . 573
Source NAT and Destination NAT . . . 574
NAT Rule Capacities . . . 575
Dynamic IP and Port NAT Oversubscription . . . 576
Dataplane NAT Memory Statistics . . . 577
Configure NAT . . . 579
LACP . . . 587
LACP Settings . . . 587
Configure LACP . . . 590

Policy . . . 593

Policy Types . . . 594
Security Policy . . . 595
Components of a Security Policy Rule . . . 596
Security Policy Best Practices . . . 599
Policy Objects . . . 600
Security Profiles . . . 601
Antivirus Profiles . . . 602
Anti-spyware Profiles . . . 603
Vulnerability Protection Profiles . . . 604
URL Filtering Profiles . . . 605
Data Filtering Profiles . . . 606
File Blocking Profiles . . . 608
DoS Protection Profile . . . 609
Zone Protection Profiles . . . 610
Security Profile Group . . . 611
Enumeration of Rules Within a Rulebase . . . 615
Use Tags to Visually Distinguish Objects . . . 617
Tips and Tricks for Searching for Objects in Security Policy . . . 619
Search Security Policy Rules for Security Profiles . . . 619
Search for Disabled Rules in the Security Policy . . . 623
Search for Log Forwarding Details in Security Policy . . . 623
Search for Log Details in Security Policy . . . 624
Search for Scheduled Rules in Security Policy . . . 625


© Palo Alto Networks, Inc.

Use a Dynamic Block List in Policy. . . . 626
View the IP Address Limit For Your Firewall Model. . . . 626
Formatting Guidelines for Dynamic Block Lists . . . . 627
Enforce Policy with a Dynamic Block List . . . . 627
View the List of IP addresses in the Dynamic Block List. . . . 628
Retrieve a Dynamic Block List from Web Server . . . . 629
Register IP Addresses and Tags Dynamically . . . . 630
Monitor Changes in the Virtual Environment . . . . 631
Enable VM Monitoring to Track Changes on the Virtual Network . . . . 632
Attributes Monitored in the AWS and VMware Environments. . . . 635
Use Dynamic Address Groups in Policy . . . . 636
CLI Commands for Dynamic IP Addresses and Tags . . . . 639
Policy-Based Forwarding . . . . 641
PBF . . . . 642
Create a Policy-Based Forwarding Rule. . . . 644
Use Case: PBF for Outbound Access with Dual ISPs . . . . 646
DoS Protection Against Flooding of New Sessions . . . . 654
DoS Protection Against Flooding of New Sessions . . . . 655
Configure DoS Protection Against Flooding of New Sessions . . . . 659
Use the CLI to End a Single Attacking Session . . . . 662
Identify Sessions That Use an Excessive Percentage of the Packet Buffer . . . . 663
Discard a Session Without a Commit . . . . 665

Virtual Systems . . . . 667

Virtual Systems Overview . . . . 668
Virtual System Components and Segmentation . . . . 668
Benefits of Virtual Systems. . . . 669
Use Cases for Virtual Systems . . . . 669
Platform Support and Licensing for Virtual Systems . . . . 670
Restrictions on Virtual Systems . . . . 670
Administrative Roles for Virtual Systems . . . . 670
Shared Objects for Virtual Systems . . . . 670
Communication Between Virtual Systems . . . . 671
Inter-VSYS Traffic That Must Leave the Firewall . . . . 672
Inter-VSYS Traffic That Remains Within the Firewall. . . . 673
Inter-VSYS Communication Uses Two Sessions . . . . 676
Shared Gateway . . . . 677
External Zones and Shared Gateway. . . . 677
Networking Considerations for a Shared Gateway . . . . 678
Configure Virtual Systems . . . . 679
Configure Inter-Virtual System Communication within the Firewall . . . . 682
Configure a Shared Gateway . . . . 683
Virtual System Functionality with Other Features . . . . 684

© Palo Alto Networks, Inc.

© Palo Alto Networks, Inc.

# Getting Started

The following sections provide detailed steps to help you deploy a new Palo Alto Networks next-generation firewall. They provide details for integrating a new firewall into your network and configuring basic security policies and threat prevention features.

After you perform the basic configuration steps required to integrate the firewall into your network, you can use the rest of the topics in this guide to help you deploy the comprehensive enterprise security platform features as necessary to address your network security needs.

- Integrate the Firewall into Your Management Network
- Create the Security Perimeter
- Enable Basic Threat Prevention Features
- Best Practices for Completing the Firewall Deployment

# Integrate the Firewall into Your Management Network

All Palo Alto Networks firewalls provide an out-of-band management port (MGT) that you can use to perform the firewall administration functions. By using the MGT port, you separate the management functions of the firewall from the data processing functions, safeguarding access to the firewall and enhancing performance. When using the web interface, you must perform all initial configuration tasks from the MGT port even if you plan to use an in-band port for managing your device going forward.

Some management tasks, such as retrieving licenses and updating the threat and application signatures on the firewall require access to the Internet. If you do not want to enable external access to your MGT port, you will need to either set up a data port to provide access to required external services or plan to manually upload updates regularly.

The following topics describe how to perform the initial configuration steps that are necessary to integrate a new firewall into the management network and deploy it in a basic security configuration.

- Determine Your Management Strategy
- Perform Initial Configuration
- Set Up Network Access for External Services
- Register the Firewall
- Activate Licenses and Subscriptions
- Manage Content Updates
- Install Software Updates

The following topics describe how to integrate a single Palo Alto Networks next-generation firewall into your network. However, for redundancy, consider deploying a pair of firewalls in a High Availability configuration.

## Determine Your Management Strategy

The Palo Alto Networks firewall can be configured and managed locally or it can be managed centrally using Panorama, the Palo Alto Networks centralized security management system. If you have six or more firewalls deployed in your network, use Panorama to achieve the following benefits:

- Reduce the complexity and administrative overhead in managing configuration, policies, software and dynamic content updates. Using device groups and templates on Panorama, you can effectively manage device specific configuration locally on a device and enforce shared policies across all devices or device groups.
- Aggregate data from all managed firewalls and gain visibility across all the traffic on your network. The Application Command Center (**ACC**) on Panorama provides a single glass pane for unified reporting across all the firewalls, allowing you to centrally analyze, investigate and report on network traffic, security incidents and administrative modifications.

 © Palo Alto Networks, Inc.

The procedures in this document describe how to manage the firewall using the local web interface. If you want to use Panorama for centralized management, after you complete the instructions in the Perform Initial Configuration section of this guide and verify that the firewall can establish a connection to Panorama. From that point on you can use Panorama to configure your firewall centrally.

## Perform Initial Configuration

By default, the firewall has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other firewall configuration tasks. You must perform these initial configuration tasks either from the MGT interface, even if you do not plan to use this interface for your firewall management, or using a direct serial connection to the console port on the device.

| **Set Up Network Access to the Firewall** | | |
|---|---|---|
| Step 1 | Gather the required information from your network administrator. | ❑ IP address for MGT port<br>❑ Netmask<br>❑ Default gateway<br>❑ DNS server address |
| Step 2 | Connect your computer to the firewall. | You can connect to the firewall in one of the following ways:<br>• Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1). Wait a few minutes for the boot-up sequence to complete; when the device is ready, the prompt changes to the name of the firewall, for example `PA-500 login`.<br>• Connect an RJ-45 Ethernet cable from your computer to the MGT port on the firewall. From a browser, go to `https://192.168.1.1`. Note that you may need to change the IP address on your computer to an address in the 192.168.1.0 network, such as 192.168.1.2, in order to access this URL. |
| Step 3 | When prompted, log in to the firewall. | You must log in using the default username and password (admin/admin). The firewall will begin to initialize. |
| Step 4 | Configure the MGT interface. | 1. Select **Device > Setup > Management** and then edit the Management Interface Settings.<br>2. Enter the **IP Address**, **Netmask**, and **Default Gateway**.<br>3. Set the **Speed** to **auto-negotiate**.<br>4. Select which management services to allow on the interface.<br> Make sure **Telnet** and **HTTP** are not selected because these services use plaintext and are not as secure as the other services.<br>5. Click **OK**. |

| Set Up Network Access to the Firewall (Continued) | | |
|---|---|---|
| Step 5 | (Optional) Configure general firewall settings. | 1. Select **Device > Setup > Management** and edit the General Settings.<br>2. Enter a **Hostname** for the firewall and enter your network **Domain** name. The domain name is just a label; it will not be used to join the domain.<br>3. Enter the **Latitude** and **Longitude** to enable accurate placement of the firewall on the world map.<br>4. Click **OK**. |
| Step 6 | Configure DNS, time and date settings.<br>You must manually configure at least one DNS server on the firewall or it will not be able to resolve hostnames; it will not use DNS server settings from another source, such as an ISP. | 1. Select **Device > Setup > Services** and edit the Services.<br>2. On the **Services** tab, enter the IP address of the **Primary DNS Server** and optionally a **Secondary DNS Server**.<br>3. To use the virtual cluster of time servers on the Internet, enter the hostname **pool.ntp.org** as the **Primary NTP Server** or add the IP address of your **Primary NTP Server** and optionally a **Secondary NTP Server**.<br>4. To authenticate time updates from an NTP server, select the **NTP** tab, enter the **NTP Server Address**, and select the **Authentication Type** for the firewall to use.<br>5. Click **OK** to save your settings. |
| Step 7 | Set a secure password for the admin account. | 1. Select **Device > Administrators.**<br>2. Select the **admin** role.<br>3. Enter the current default password and the new password.<br>4. Click **OK** to save your settings. |
| Step 8 | Commit your changes.<br>When the configuration changes are saved, you will lose connectivity to the web interface because the IP address will have changed. | Click **Commit**. The device may take up to 90 seconds to save your changes.<br> |
| Step 9 | Connect the firewall to your network. | 1. Disconnect the firewall from your computer.<br>2. Connect the MGT port to a switch port on your management network using an RJ-45 Ethernet cable. Make sure that the switch port you cable the firewall to is configured for auto-negotiation. |
| Step 10 | Open an SSH management session to the firewall. | Using a terminal emulation software, such as PuTTY, launch an SSH session to the firewall using the new IP address you assigned to it. |

 © Palo Alto Networks, Inc.

| Set Up Network Access to the Firewall (Continued) | | |
|---|---|---|
| Step 11 | Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server, in one of the following ways:<br>• If you do not want to allow external network access to the MGT interface, you will need to set up a data port to retrieve required service updates. Continue to Set Up Network Access for External Services.<br>• If you do plan to allow external network access to the MGT interface, verify that you have connectivity and then proceed to Register the Firewall and Activate Licenses and Subscriptions. | If you cabled your MGT port for external network access, verify that you have access to and from the firewall by using the ping utility from the CLI. Make sure you have connectivity to the default gateway, DNS server, and the Palo Alto Networks Update Server as shown in the following example:<br>`admin@PA-200> ping host updates.paloaltonetworks.com`<br>`PING updates.paloaltonetworks.com (67.192.236.252) 56(84) bytes of data.`<br>`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=40.5 ms`<br>`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=53.6 ms`<br>`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=79.5 ms`<br>After you have verified connectivity, press Ctrl+C to stop the pings. |

## Set Up Network Access for External Services

By default, the firewall uses the MGT interface to access remote services, such as DNS servers, content updates, and license retrieval. If you do not want to enable external network access to your management network, you must set up a data port to provide access to these required external services.

This task requires familiarity with firewall interfaces, zones, and policies. For more information on these topics, see Create the Security Perimeter.

| Set Up a Data Port for Access to External Services | | |
|---|---|---|
| Step 1 | Decide which port you want to use for access to external services and connect it to your switch or router port. | The interface you use will need to have a static IP address. |
| Step 2 | Log in to the web interface. | Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page. |

| Set Up a Data Port for Access to External Services (Continued) | | |
|---|---|---|
| Step 3 | (Optional) The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and zones). If you do not plan to use this virtual wire configuration, you must manually delete the configuration to prevent it from interfering with other interface settings you define. | You must delete the configuration in the following order:<br>1. To delete the default security policy, select **Policies > Security**, select the rule, and click **Delete**.<br>2. Next, delete the default virtual wire by selecting **Network > Virtual Wires**, selecting the virtual wire and clicking **Delete**.<br>3. To delete the default trust and untrust zones, select **Network > Zones**, select each zone and click **Delete**.<br>4. Finally, delete the interface configurations by selecting **Network > Interfaces** and then select each interface (ethernet1/1 and ethernet1/2) and click **Delete**.<br>5. **Commit** the changes. |
| Step 4 | Configure the interface. | 1. Select **Network > Interfaces** and select the interface that corresponds to the port you cabled in Step 1.<br>2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**.<br>4. In the Zone dialog, define a **Name** for new zone, for example L3-trust, and then click **OK**.<br>5. Select the **IPv4** tab, select the **Static** radio button, and click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.254/24.<br>6. Select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**.<br>7. Enter a **Name** for the profile, such as allow_ping, and then select the services you want to allow on the interface. These services provide management access to the device, so only select the services that correspond to the management activities you want to allow on this interface. For example, if you plan to use the MGT interface for device configuration tasks through the web interface or CLI, you would not want to enable HTTP, HTTPS, SSH, or Telnet so that you could prevent unauthorized access through this interface. For the purposes of allowing access to the external services you probably only need to enable **Ping** and then click **OK**.<br>8. To save the interface configuration, click **OK**. |

 © Palo Alto Networks, Inc.

**Set Up a Data Port for Access to External Services (Continued)**

| | |
|---|---|
| Step 5 Because the firewall uses the MGT interface by default to access the external services it requires, you must change the interface the firewall uses to send these requests by editing the service routes. | 1. Select **Device > Setup > Services > Service Route Configuration**.<br> 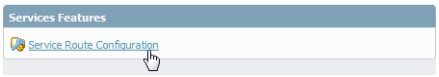 <br>For the purposes of activating your licenses and getting the most recent content and software updates, you will want to change the service route for **DNS**, **Palo Alto Updates**, **URL Updates**, and **WildFire**.<br>2. Click the **Customize** radio button, and select one of the following:<br>• For a predefined service, select **IPv4** or **IPv6** and click the link for the service for which you want to modify the **Source Interface** and select the interface you just configured.<br>If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.<br>• To create a service route for a custom destination, select **Destination**, and click **Add**. Enter a **Destination** name and select a **Source Interface**. If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.<br>3. Click **OK** to save the settings.<br>4. Repeat steps 2-3 above for each service route you want to modify.<br>5. **Commit** your changes. |

**Set Up a Data Port for Access to External Services (Continued)**

| | | |
|---|---|---|
| Step 6 | Configure an external-facing interface and an associated zone and then create security and NAT policy rules to allow the firewall to send service requests from the internal zone to the external zone. | 1. Select **Network > Interfaces** and then select your external-facing interface. Select **Layer3** as the **Interface Type**, **Add** the **IP** address (on the **IPv4** or **IPv6** tab), and create the associated **Security Zone** (on the **Config** tab), such as l3-untrust. You do not need to set up management services on this interface.<br>2. To set up a security rule that allows traffic from your internal network to the Palo Alto Networks update server and external DNS servers, select **Policies > Security** and click **Add**. For the purposes of initial configuration, you can create a simple rule that allows all traffic from l3-trust to l3-untrust as follows:<br> <br>3. If you are using a private IP address on the internal-facing interface, you will need to create a Source NAT rule to translate the address to a publicly routable address. Select **Policies > NAT** and then click **Add**. At a minimum you must define a name for the rule (**General** tab), specify a source and destination zone, l3-trust to l3-untrust in this case (**Original Packet** tab), and define the source address translation settings (**Translated Packet** tab) and then click **OK**.<br> <br>4. **Commit** your changes. |
| Step 7 | Verify that you have connectivity from the data port to the external services, including the default gateway, DNS server, and the Palo Alto Networks Update Server.<br><br>After you verify you have the required network connectivity, continue to Register the Firewall and Activate Licenses and Subscriptions. | Launch the CLI and use the ping utility to verify that you have connectivity. Keep in mind that by default pings are sent from the MGT interface, so in this case you must specify the source interface for the ping requests as follows:<br>`admin@PA-200> ping source 192.168.1.254 host updates.paloaltonetworks.com`<br>`PING updates.paloaltonetworks.com (67.192.236.252) from 192.168.1.254 : 56(84) bytes of data.`<br>`64 bytes from 67.192.236.252: icmp_seq=1 ttl=242 time=56.7 ms`<br>`64 bytes from 67.192.236.252: icmp_seq=2 ttl=242 time=47.7 ms`<br>`64 bytes from 67.192.236.252: icmp_seq=3 ttl=242 time=47.6 ms`<br>`^C`<br>After you have verified connectivity, press Ctrl+C to stop the pings. |

 © Palo Alto Networks, Inc.

## Register the Firewall

| Register the Firewall | | |
|---|---|---|
| Step 1 | Log in to the web interface. | Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page. |
| Step 2 | Locate your serial number and copy it to the clipboard. | On the **Dashboard**, locate your **Serial Number** in the General Information section of the screen. |
| Step 3 | Go to the Palo Alto Networks Support site. | In a new browser tab or window, go to https://support.paloaltonetworks.com. |
| Step 4 | Register the device. The way you register depends on whether you already have a login to the support site. | • If this is the first Palo Alto Networks device you are registering and you do not yet have a login, click **Register** on the right side of the page. To register, you must provide your sales order number or customer ID, and the serial number of your firewall (which you can paste from your clipboard) or the authorization code you received with your order. You will also be prompted to set up a username and password for access to the Palo Alto Networks support community.<br> <br>• If you already have a support account, log in and then click **My Devices**. Scroll down to Register Device section at the bottom of the screen and enter the serial number of your firewall (which you can paste from your clipboard), your city and postal code and then click **Register Device**. |

## Activate Licenses and Subscriptions

Before you can start using your firewall to secure the traffic on your network, you must activate the licenses for each of the services you purchased. Available licenses and subscriptions include the following:

- Threat Prevention—Provides antivirus, anti-spyware, and vulnerability protection.
- Decryption Port Mirroring—Provides the ability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as NetWitness or Solera—for archiving and analysis.
- URL Filtering—In order to create policy rules based on dynamic URL categories, you must purchase and install a subscription for one of the supported URL filtering databases: PAN-DB or BrightCloud. For more information about URL filtering, see Control Access to Web Content.

© Palo Alto Networks, Inc.

- Virtual Systems—This license is required to enable support for multiple virtual systems on PA-2000 and PA-3000 Series firewalls. In addition, you must purchase a Virtual Systems license if you want to increase the number of virtual systems beyond the base number provided by default on PA-4000 Series, PA-5000 Series, and PA-7050 firewalls (the base number varies by platform). The PA-500, PA-200, and VM-Series firewalls do not support virtual systems.
- WildFire—Although basic WildFire support is included as part of the Threat Prevention license, the WildFire subscription service provides enhanced services for organizations that require immediate coverage for threats, enabling sub-hourly WildFire signature updates, advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet), as well as the ability to upload files using the WildFire API. A WildFire subscription is also required if your firewalls will be forwarding files to a private WF-500 WildFire appliance.
- GlobalProtect—Provides mobility solutions and/or large-scale VPN capabilities. By default, you can deploy a single GlobalProtect portal and gateway (without HIP checks) without a license. However, if you want to deploy multiple gateways, you must purchase a portal license (one-time, permanent license). If you want to use host checks you will also need gateway licenses (subscription) for each gateway.

| **Activate Licenses** | | |
|---|---|---|
| Step 1 | Locate the activation codes for the licenses you purchased. | When you purchased your subscriptions you should have received an email from Palo Alto Networks customer service listing the activation code associated with each subscription. If you cannot locate this email, contact customer support to obtain your activation codes before you proceed. |
| Step 2 | Launch the web interface and go to the license page. | Select **Device > Licenses**. |
| Step 3 | Activate each license you purchased. | After purchasing your licenses/subscriptions activate them in one of the following ways:<br>• **Retrieve license keys from license server**—Use this option if you activated your license on the support portal.<br>• **Activate feature using authorization code**—Use this option to enable purchased subscriptions using an authorization code for licenses that have not been previously activated on the support portal. When prompted, enter the **Authorization Code** and then click **OK**.<br>• **Manually upload license key**—Use this option if your device does not have connectivity to the Palo Alto Networks support site. In this case, you must download a license key file from the support site on an Internet connected computer and then upload to the device. |
| Step 4 | Verify that the license was successfully activated | On the **Device > Licenses** page, verify that the license was successfully activated. For example, after activating the WildFire license, you should see that the license is valid:<br> |

 © Palo Alto Networks, Inc.

| Activate Licenses (Continued) | |
|---|---|
| Step 5 (WildFire subscriptions only) Perform a commit to complete WildFire subscription activation. | After activating a WildFire subscription, a commit is required for the firewall to begin forwarding advanced file types:<br>• Commit any pending changes.<br>• Make a minor change and perform a commit. For example, update a rule description and commit the change.<br> Check the file blocking profile rule used for WildFire forwarding includes the advanced file types that are now supported with the WildFire subscription. If the rule is already set to forward **Any** file type, or no change to the rule is required, make a minor edit to the rule description and perform a commit. |

## Manage Content Updates

In order to stay ahead of the changing threat and application landscape, Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to the Palo Alto Networks devices. The devices access the web resources in the CDN to perform various App-ID and Content-ID functions. By default, the devices use the management port to access the CDN infrastructure for application updates, threat and antivirus signature updates, BrightCloud and PAN-DB database updates and lookups, and access to the Palo Alto Networks WildFire Cloud. To ensure that you are always protected from the latest threats (including those that have not yet been discovered), you must ensure that you keep your devices up-to-date with the latest updates published by Palo Alto Networks.

The following content updates are available, depending on which subscriptions you have:

Although you can manually download and install content updates at any time, as a best practice you should Schedule each update. Scheduled updates occur automatically.

- **Antivirus**—Includes new and updated antivirus signatures, including signatures discovered by the WildFire cloud service. You must have a Threat Prevention subscription to get these updates. New antivirus signatures are published daily.
- **Applications**—Includes new and updated application signatures. This update does not require any additional subscriptions, but it does require a valid maintenance/support contract. New application updates are published weekly.
- **Applications and Threats**—Includes new and updated application and threat signatures. This update is available if you have a Threat Prevention subscription (and you get it instead of the Applications update). New Applications and Threats updates are published weekly.
- **GlobalProtect Data File**—Contains the vendor-specific information for defining and evaluating host information profile (HIP) data returned by GlobalProtect agents. You must have a GlobalProtect portal and GlobalProtect gateway license in order to receive these updates. In addition, you must create a schedule for these updates before GlobalProtect will function.

- **BrightCloud URL Filtering**—Provides updates to the BrightCloud URL Filtering database only. You must have a BrightCloud subscription to get these updates. New BrightCloud URL database updates are published daily. If you have a PAN-DB license, scheduled updates are not required as devices remain in-sync with the servers automatically.
- **WildFire**—Provides near real-time malware and antivirus signatures created as a result of the analysis done by the WildFire cloud service. Without the subscription, you must wait 24 to 48 hours for the signatures to roll into the Applications and Threat update.

If your firewall does not have Internet access from the management port, you can download content updates from the Palo Alto Networks Support portal and then Upload them to your firewall.

If your firewall is deployed behind existing firewalls or proxy servers, access to these external resources might be restricted using access control lists that allow the firewall to only access a hostname or an IP address. In such cases, to allow access to the CDN, set the update server address to use the hostname `staticupdates.paloaltonetworks.com` or the IP address `199.167.52.15`.

| **Download the Latest Databases** | |
|---|---|
| Step 1 Verify that the firewall points to the CDN infrastructure. | Select **Device > Setup > Services**.<br>• As a best practice, set the **Update Server** to access `updates.paloaltonetworks.com`. This allows the firewall to receive content updates from the server to which it is closest in the CDN infrastructure.<br>• (Optional) If the firewall has restricted access to the Internet, set the update server address to use the hostname `staticupdates.paloaltonetworks.com` or the IP address `199.167.52.15`.<br>• For additional security, select **Verify Update Server Identity**. The firewall verifies that the server from which the software or content package is download has an SSL certificate signed by a trusted authority. |
| Step 2 Launch the web interface and go to the Dynamic Updates page. | Select **Device > Dynamic Updates**. |

 © Palo Alto Networks, Inc.

**Download the Latest Databases (Continued)**

Step 3 Check for the latest updates.

Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available:

- **Download**—Indicates that a new update file is available. Click the link to begin downloading the file directly to the firewall. After successful download, the link in the **Action** column changes from **Download** to **Install**.

You cannot download the antivirus database until you have installed the Application and Threats database.

- **Upgrade**—Indicates that there is a new version of the BrightCloud database available. Click the link to begin the download and installation of the database. The database upgrade begins in the background; when completed a check mark displays in the **Currently Installed** column. Note that if you are using PAN-DB as your URL filtering database you will not see an upgrade link because the PAN-DB database automatically stays in sync with the server.

To check the status of an action, click **Tasks** (on the lower right-hand corner of the window).

- **Revert**—Indicates that the corresponding software version has been downloaded previously. You can choose to revert to the previously installed version of the update.

Step 4 Install the updates.

Installation can take up to 20 minutes on a PA-200, PA-500, or PA-2000 device and up to two minutes on a PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7050, or VM-Series firewall.

Click the **Install** link in the **Action** column. When the installation completes, a check mark displays in the **Currently Installed** column.

| Download the Latest Databases (Continued) | | |
|---|---|---|
| Step 5 | Schedule each update.<br><br>Repeat this step for each update you want to schedule.<br><br>Stagger the update schedules because the firewall can only download one update at a time. If you schedule the updates to download during the same time interval, only the first download will succeed. | 1. Set the schedule of each update type by clicking the **None** link.<br> <br>2. Specify how often you want the updates to occur by selecting a value from the **Recurrence** drop-down. The available values vary by content type (WildFire updates are available **Every 15 minutes**, **Every 30 minutes** or **Every Hour** whereas all other content types can be scheduled for **Daily** or **Weekly** update).<br>3. Specify the **Time** and (or, minutes past the hour in the case of WildFire), if applicable depending on the **Recurrence** value you selected, **Day** of the week that you want the updates to occur.<br>4. Specify whether you want the system to **Download Only** or, as a best practice, **Download And Install** the update.<br>5. In rare instances, errors in content updates may be found. For this reason, you may want to delay installing new updates until they have been released for a certain number of hours. You can specify how long after a release to wait before performing a content update by entering the number of hours to wait in the **Threshold (Hours)** field.<br>6. Click **OK** to save the schedule settings.<br>7. Click **Commit** to save the settings to the running configuration. |

## Install Software Updates

When installing a new firewall, it is a good idea to upgrade to the latest software update (or to the update version recommended by your reseller or Palo Alto Networks Systems Engineer) to take advantage of the latest fixes and security enhancements. Note that before updating the software, you should first make sure you have the latest content updates as detailed in the previous section (the Release Notes for a software update specify the minimum Content Release version supported in the release).

| Update PAN-OS | | |
|---|---|---|
| Step 1 | Launch the web interface and go to the Software page. | Select **Device > Software**. |
| Step 2 | Check for software updates. | Click **Check Now** to check for the latest updates. If the value in the **Action** column is **Download** it indicates that an update is available. |

 © Palo Alto Networks, Inc.

| Update PAN-OS (Continued) | |
|---|---|
| Step 3 Download the update.<br><br>If your firewall does not have Internet access from the management port, you can download the software update from the Palo Alto Networks Support portal. You can then manually **Upload** them to your firewall. | Locate the version you want and then click **Download**. When the download completes, the value in the Action column changes to **Install**.<br> 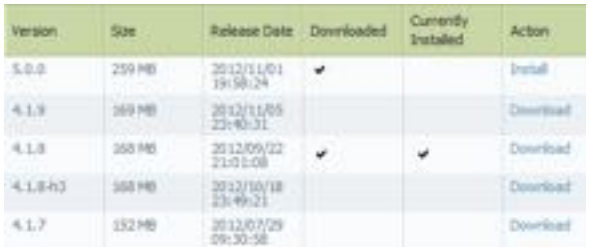 <br> |
| Step 4 Install the update. | 1. Click **Install**.<br>2. Reboot the firewall:<br>• If you are prompted to reboot, click **Yes**.<br> <br><br>• If you are not prompted to reboot, select **Device > Setup > Operations** and click **Reboot Device** in the Device Operations section of the screen.<br> <br> |

# Create the Security Perimeter

Traffic must pass through the firewall in order for the firewall to manage and control it. Physically, traffic enters and exits the firewall through *interfaces*. The firewall determines how to act on a packet based on whether the packet matches a *security policy rule*. At the most basic level, each security policy rule must identify where the traffic came from and where it is going. On a Palo Alto Networks next-generation firewall, security policy rules are applied between zones. A *zone* is a grouping of interfaces (physical or virtual) that provides an abstraction for an area of trust for simplified policy enforcement. For example, in the following topology diagram, there are three zones: Trust, Untrust, and DMZ. Traffic can flow freely within a zone, but traffic will not be able to flow between zones until you define a security policy rule that allows it.

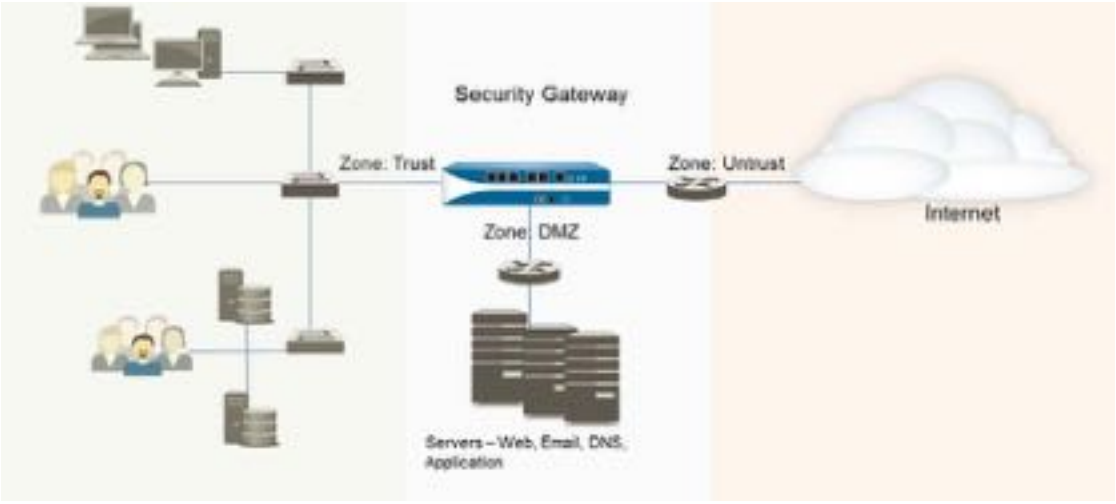


The following topics describe the components of the security perimeter and provide steps for configuring the firewall interfaces, defining zones, and setting up a basic security policy that allows traffic from your internal zone to the Internet and to the DMZ. By initially creating a basic security policy rulebase like this, you will be able to analyze the traffic running through your network and use this information to define more granular policies for safely enabling applications while preventing threats.

- Basic Interface Deployments
- About Security Policy
- Plan the Deployment
- Configure Interfaces and Zones
- Set Up Basic Security Policies

If you use private IP addresses in your internal networks, you will also need to configure network address translation (NAT); see Networking for NAT concepts and configuration tasks.

## Basic Interface Deployments

All Palo Alto Networks next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, enabling you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from

 © Palo Alto Networks, Inc.

virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports. The following sections provide basic information on each type of deployment.

- Virtual Wire Deployments
- Layer 2 Deployments
- Layer 3 Deployments

For more detailed deployment information, refer to Designing Networks with Palo Alto Networks Firewalls.

## Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together. By using a virtual wire, you can install the firewall in any network environment without reconfiguring adjacent devices. If necessary, a virtual wire can block or allow traffic based on the virtual LAN (VLAN) tag values. You can also create multiple subinterfaces and classify traffic according to an IP Address (address, range, or subnet), VLAN, or a combination of the two.

By default, the virtual wire (named **default-vwire**) binds Ethernet ports 1 and 2 and allows all untagged traffic. Choose this deployment to simplify installation and configuration and/or avoid configuration changes to surrounding network devices.

A virtual wire is the default configuration, and should be used only when no switching or routing is needed. If you do not plan to use the default virtual wire, you should manually delete the configuration before proceeding with interface configuration to prevent it from interfering with other interface settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up a Data Port for Access to External Services.

## Layer 2 Deployments

In a Layer 2 deployment, the firewall provides switching between two or more interfaces. Each group of interfaces must be assigned to a VLAN object in order for the firewall to switch between them. The firewall will perform VLAN tag switching when Layer 2 subinterfaces are attached to a common VLAN object. Choose this option when switching is required.

For more information on Layer 2 deployments, refer to the Layer 2 Networking Tech Note and/or the Securing Inter VLAN Traffic Tech Note.

## Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between ports. An IP address must be assigned to each interface and a virtual router must be defined to route the traffic. Choose this option when routing is required.

You must assign an IP address to each physical Layer 3 interface you configure. You can also create logical subinterfaces for each physical Layer 3 interface that allows you to segregate the traffic on the interface based on VLAN tag (when VLAN trunking is in use) or by IP address, for example for multi-tenancy.

In addition, because the firewall must route traffic in a Layer 3 deployment, you must configure a virtual router. You can configure the virtual router to participate with dynamic routing protocols (BGP, OSPF, or RIP) as well as adding static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that are not shared between virtual routers, enabling you to configure different routing behaviors for different interfaces.

The configuration example in this chapter illustrates how to integrate the firewall into your Layer 3 network using static routes. For information on other types of routing integrations, refer to the following documents:

- How to Configure OSPF Tech Note
- How to Configure BGP Tech Note

## About Security Policy

Security Policy protects network assets from threats and disruptions and aids in optimally allocating network resources for enhancing productivity and efficiency in business processes. On the Palo Alto Networks firewall, security policy rules determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service.

For traffic that doesn't match any defined rules, the default rules apply. The default rules—displayed at the bottom of the security rulebase—are predefined to allow all intrazone (within the zone) traffic and deny all interzone (between zones) traffic. Although these rules are part of the pre-defined configuration and are read-only by default, you can override them and change a limited number of settings, including the tags, action (allow or deny), log settings, and security profiles.

Security policies rules are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log, if logging is enabled for that rule. The logging options are configurable for each rule, and can for example be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.

▲ About Policy Objects
▲ About Security Profiles

### About Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With Policy Objects that are a collective unit, you can reference the object in security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

Some examples of address and application policy objects are shown in the security policies that are included in Create Security Rules. For information on the other policy objects, see Enable Basic Threat Prevention Features.

 © Palo Alto Networks, Inc.

## About Security Profiles

While security policies enable you to allow or deny traffic on your network, security profiles help you define an *allow but scan* rule, which scan allowed applications for threats. When traffic matches the allow rule defined in the security policy, the Security Profiles that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.

Security profiles are not used in the match criteria of a traffic flow. The security profile is applied to scan traffic after the application or category is allowed by the security policy.

The different types of security profiles that can be attached to security policies are: Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking, and Data Filtering. The firewall provides default security profiles that you can use out of the box to begin protecting your network from threats. See Create Security Rules for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, you can create custom profiles. See Scan Traffic for Threats for more information.

# Plan the Deployment

Before you begin configuring interfaces and zones, take some time to plan the zones you will need based on the different usage requirements within your organization. In addition, you should gather all of the configuration information you will need ahead of time. At a basic level, you should plan which interfaces will belong to which zones. For Layer 3 deployments you'll also need to obtain the required IP addresses and network configuration information from your network administrator, including information on how to configure the routing protocol or static routes required for the virtual router configuration. The example in this chapter will be based on the following topology:

**Figure: Layer 3 Topology Example**

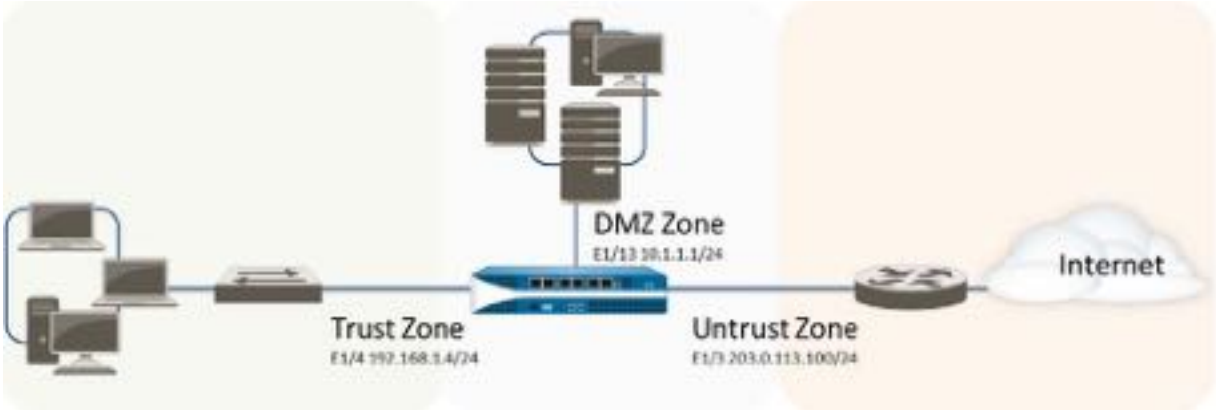


The following table shows the information we will use to configure the Layer 3 interfaces and their corresponding zones as shown in the sample topology.

| Zone | Deployment Type | Interface(s) | Configuration Settings |
|---|---|---|---|
| Untrust | L3 | Ethernet1/3 | **IP address:** 203.0.113.100/24<br>**Virtual router:** default<br>**Default route:** 0.0.0.0/0<br>**Next hop:** 203.0.113.1 |
| Trust | L3 | Ethernet1/4 | **IP address:** 192.168.1.4/24<br>**Virtual router:** default |
| DMZ | L3 | Ethernet1/13 | **IP address:** 10.1.1.1/24<br>**Virtual router:** default |

## Configure Interfaces and Zones

After you plan your zones and the corresponding interfaces, you can configure them on the device. The way you configure each interface depends on your network topology.

The following procedure shows how to configure a Layer 3 deployment as depicted in Figure: Layer 3 Topology Example.

The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and virtual router). If you do not plan to use the default virtual wire, you must manually delete the configuration and commit the change before proceeding to prevent it from interfering with other settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up a Data Port for Access to External Services.

| Set Up Interfaces and Zones | |
|---|---|
| Step 1 Configure a default route to your Internet router. | 1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog.<br>2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and enter the route in the **Destination** field (for example, 0.0.0.0/0).<br>3. Select the **IP Address** radio button in the **Next Hop** field and then enter the IP address and netmask for your Internet gateway (for example, 203.00.113.1).<br>4. Click **OK** twice to save the virtual router configuration. |

 © Palo Alto Networks, Inc.

| Set Up Interfaces and Zones (Continued) | | |
|---|---|---|
| Step 2 | Configure the external interface (the interface that connects to the Internet). | 1. Select **Network > Interfaces** and then select the interface you want to configure. In this example, we are configuring Ethernet1/3 as the external interface.<br>2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.<br>3. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Untrust, and then click **OK**.<br>4. In the **Virtual Router** drop-down, select **default**.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 208.80.56.100/24.<br>6. To enable you to ping the interface, select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. Enter a **Name** for the profile, select **Ping** and then click **OK**.<br>7. To save the interface configuration, click **OK**. |
| Step 3 | Configure the interface that connects to your internal network.<br> In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT. | 1. Select **Network > Interfaces** and select the interface you want to configure. In this example, we are configuring Ethernet1/4 as the internal interface.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Trust, and then click **OK**.<br>4. Select the same Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |

| Set Up Interfaces and Zones (Continued) | | |
|---|---|---|
| Step 4 | Configure the interface that connects to the DMZ. | 1. Select the interface you want to configure.<br>2. Select **Layer3** from the **Interface Type** drop-down. In this example, we are configuring Ethernet1/13 as the DMZ interface.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example DMZ, and then click **OK**.<br>4. Select the Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 5 | Save the interface configuration. | Click **Commit**. |
| Step 6 | Cable the firewall. | Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment. |
| Step 7 | Verify that the interfaces are active. | From the web interface, select **Network > Interfaces** and verify that icon in the Link State column is green. You can also monitor link state from the **Interfaces** widget on the **Dashboard**.<br> 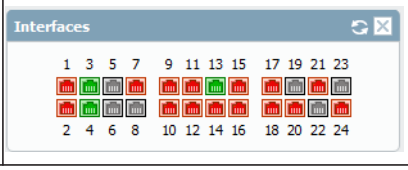 |

## Set Up Basic Security Policies

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Captive Portal, Denial of Service, and Zone protection policies. All these different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. This section covers basic security policies and the default security profiles:

- Create Security Rules
- Test Your Security Policies
- Monitor the Traffic on Your Network

 © Palo Alto Networks, Inc.

## Create Security Rules

Security policies reference security zones and enable you to allow, restrict, and track traffic on your network. Because each zone implies a level of trust, the implicit rule for passing traffic between two different zones is deny, and the traffic within a zone is permitted. To allow traffic between two different zones, you must create a security rule that allows traffic to flow between them.

While setting up the basic framework for securing the enterprise perimeter, it's good idea to start with a simple security policy that allows traffic between the different zones without being too restrictive. As illustrated in the following section, our objective is to minimize the likelihood of breaking applications that users on the network need access to, while providing visibility into the applications and the potential threats for your network.

When defining policies make sure that you do not create a policy that denies all traffic from *any* source zone to *any* destination zone as this will break intra-zone traffic that is implicitly allowed. By default, intra-zone traffic is permitted because the source and destination zones are the same and therefore share the same level of trust.

## Define Basic Security Rules

**Step 1** Permit Internet access for all users on the enterprise network.

Zone: Trust to Untrust

By default, the firewall includes a security rule named *rule1* that allows all traffic from Trust zone to Untrust zone. You can either delete the rule or modify the rule to reflect your zone-naming convention.

To safely enable applications that are required for day-to-day business operations we will create a simple rule that allows access to the Internet. To provide basic threat protection, we will attach the default security profiles available on the firewall.

1. Select **Policies > Security** and click **Add.**
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, Set the **Destination Zone** to Untrust.

   To scan policy rules and visually identify the zones on each rule, create a tag with the same name as the zone. For example, to color code the Trust zone as green, select **Objects > Tags,** click **Add** and **Name** the tag Trust, and select the **Color** green.

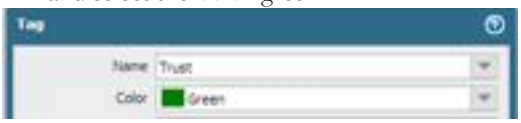


5. In the **Service/ URL Category** tab, select service-http and service-https.
6. In the **Actions** tab, complete these tasks:
   a. Set the **Action Setting** to **Allow**.
   b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection and URL filtering, under **Profile Setting.**
7. Verify that logging is enabled at the end of a session under **Options**. Only traffic that matches a security rule will be logged.

**Step 2** Permit users on the internal network to access the servers in the DMZ.

Zone: Trust to DMZ

If using IP addresses for configuring access to the servers in the DMZ, always refer to the original IP addresses in the packet (i.e. the pre-NAT addresses), and the post-NAT zone.

1. Click **Add** in the **Policies > Security** section.
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, set the **Destination Zone** to DMZ.
5. In the **Service/ URL Category** tab, make sure the **Service** is set to **application-default**.
6. In the **Actions** tab, set the **Action Setting** to Allow.
7. Leave all the other options at the default values.

 © Palo Alto Networks, Inc.

**Define Basic Security Rules (Continued)**

Step 3 Restrict access from the Internet to the servers on the DMZ to specific server IP addresses only.

For example, you might only allow users to access the webmail servers from outside.

Zone: Untrust to DMZ

To restrict inbound access to the DMZ from the Internet, configure a rule that allows access only to specific servers IP addresses and on the default ports that the applications use.

1. Click **Add** to add a new rule, and give it a descriptive name.
2. In the **Source** tab, set the **Source Zone** to Untrust.
3. In the **Destination** tab, set the **Destination Zone** to DMZ.
4. Set the **Destination Address** to the **Public web server** address object you created earlier. The public web server address object references the public IP address—208.80.56.11/24—of the web server that is accessible on the DMZ.
5. Select the webmail application in the **Application** tab.

   The **Service** is set to **application-default** by default.

6. Set the **Action Setting** to **Allow**.

Step 4 Allow access from the DMZ to your internal network (Trust zone). To minimize risk, you will allow traffic only between specific servers and destination addresses. For example, if you have an application server on the DMZ that needs to communicate with a specific database server in your Trust zone, create a rule to allow traffic between a specific source to a specific destination.

Zone: DMZ to Trust

1. Click **Add** to add a new rule, and give it a descriptive name.
2. Set the **Source Zone** to DMZ.
3. Set the **Destination Zone** to Trust.
4. Create a an address object that specifies the server(s) on your Trust zone that can be accessed from the DMZ.

5. In the **Destination** tab on the Security Policy rule, set the **Destination Address** to the Address object you created above.
6. In the **Actions** tab, complete these tasks:
   a. Set the **Action Setting** to **Allow**.
   b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection, under **Profile Setting.**
   c. In the Other Settings section, select the option to **Disable Server Response Inspection**. This setting disables the antivirus and anti-spyware scanning on the server-side responses, and thus reduces the load on the firewall.

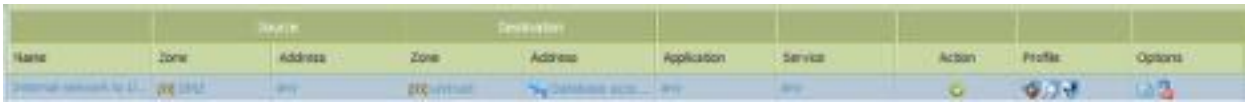

| Define Basic Security Rules (Continued) | |
|---|---|
| Step 5 Enable the servers on the DMZ to obtain updates and hot fixes from the Internet. Say, for example, you would like to allow the Microsoft Update service.<br><br>Zone: DMZ to Untrust | 1. Add a new rule and give it a descriptive label.<br>2. Set the **Source Zone** to DMZ.<br>3. Set the **Destination Zone** to Untrust.<br>4. Create an application group to specify the applications that you would like to allow. In this example, we allow Microsoft updates (ms-updates) and dns.<br><br> 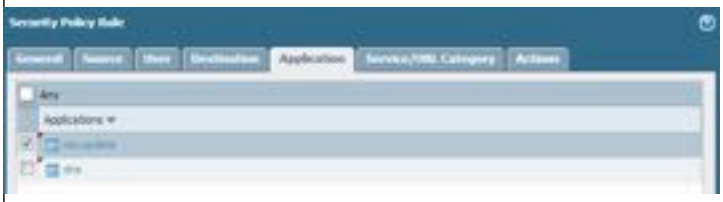 <br><br>The **Service** is set to **application-default** by default. This allows the firewall to permit the applications only when they use the standard ports associated with these applications.<br>5. Set the **Action Setting** to **Allow**.<br>6. Attach the default profiles for antivirus, anti-spyware, and vulnerability protection, under **Profiles**. |

| | |
|---|---|
| Step 6 Save your policies to the running configuration on the device. | Click **Commit**. |

 © Palo Alto Networks, Inc.

## Test Your Security Policies

To verify that you have set up your basic policies effectively, test whether your security policies are being evaluated and determine which security rule applies to a traffic flow.

| **Verify Policy Match Against a Flow** | |
|---|---|
| To verify the policy rule that matches a flow, use the following CLI command:<br>**`test security-policy-match source`** `<IP_address>` **`destination`** `<IP_address>` **`destination port`** `<port_number>` **`protocol`** `<protocol_number>`<br>The output displays the best rule that matches the source and destination IP address specified in the CLI command. | For example, to verify the policy rule that will be applied for a server on the DMZ with the IP address 208.90.56.11 when it accesses the Microsoft update server, you will try the following command:<br>`test security-policy-match source 208.80.56.11 destination 176.9.45.70 destination-port 80 protocol 6`<br><br>`"Updates-DMZ to Internet" {`<br>`from dmz;`<br>`source any;`<br>`source-region any;`<br>`to untrust;`<br>`destination any;`<br>`destination-region any;`<br>`user any;`<br>`category any;`<br>`application/service[ dns/tcp/any/53 dns/udp/any/53 dns/udp/any/5353 ms-update/tcp/any/80 ms-update/tcp/any/443];`<br>`action allow;`<br>`terminal yes;` |

## Monitor the Traffic on Your Network

Now that you have a basic security policy in place, you can review the statistics and data in the Application Command Center (ACC), traffic logs, and the threat logs to observe trends on your network, to identify where you need to create more granular policies.

Unlike traditional firewalls that use port or protocol to identify applications, the Palo Alto Networks firewalls use the application signature (the App-ID technology) to monitor applications. The application signature is based on unique application properties and related transaction characteristics in combination with the port or protocol. Therefore, even when the traffic uses the right port/protocol, the firewall can deny access to content because the application signature is not a match. This feature allows you to safely enable applications by allowing parts of the application while blocking or controlling functions within the same application. For example, if you allow the application *web-browsing* a user will be able to access content on the Internet. Then, if a user goes to Facebook and then goes on to play Scrabble on Facebook, the firewall will identify the application shifts and recognize Facebook as an *application* and Scrabble as a *Facebook-app*. Therefore, if you create a specific rule that blocks Facebook applications, the user will be denied access to Scrabble while still being able to access Facebook.

| Monitor Network Traffic | |
|---|---|
| • Use the Application Command Center. | In the ACC, review the most used applications and the high-risk applications on your network. The ACC graphically summarizes the log information to highlight the applications traversing the network, who is using them (with User-ID enabled), and the potential security impact of the content to help you identify what is happening on the network in real time. You can then use this information to create appropriate security policies that block unwanted applications, while allowing and enabling applications in a secure manner. |
| • Determine what updates/modifications are required for your network security rules and implement the changes. | For example:<br>• Evaluate whether to allow content based on schedule, users, or groups<br>• Allow or control certain applications or functions within an application<br>• Decrypt and inspect content<br>• Allow but scan for threats and exploits<br>For information on refining your security policies and for attaching custom security profiles, see Enable Basic Threat Prevention Features. |
| • View the Log Files. | Specifically, view the traffic and threat logs (**Monitor > Logs**).<br>Traffic logs are dependent on how your security policies are defined and setup to log traffic. The ACC tab, however, records applications and statistics regardless of policy configuration; it shows all traffic that is allowed on your network, therefore it includes the inter zone traffic that is allowed by policy and the same zone traffic that is allowed implicitly |
| • Interpret the URL Filtering Logs | Review the URL filtering logs to scan through alerts, denied categories/URL. URL logs are generated when a traffic matches a security rule that has a URL filtering profile attached with an action of alert, continue, override or block. |

 © Palo Alto Networks, Inc.

# Enable Basic Threat Prevention Features

The Palo Alto Networks next-generation firewall has unique threat prevention capabilities that allow it to protect your network from attack despite evasive, tunneled, or circumvention techniques. The threat prevention features on the firewall include the WildFire service, the Security Profiles that support Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking and Data Filtering capabilities and the Denial of Service (DoS) and Zone protection functionality.

Before you can apply threat prevention features, you must first configure zones—to identify one or more source or destination interfaces—and security policies. To configure interfaces, zones, and the policies that are needed to apply threat prevention features, see Configure Interfaces and Zones and Set Up Basic Security Policies.

To begin protecting your network from threats start here:

- Enable WildFire
- Scan Traffic for Threats
- Control Access to Web Content

## Enable WildFire

The WildFire service is included as part of the base product. The WildFire service enables the firewall to forward attachments to a sandbox environment where applications are run to detect any malicious activity. As new malware is detected by the WildFire system, malware signatures are automatically generated and are made available within 24-48 hours in the antivirus daily downloads. Your threat prevention subscription entitles you for antivirus signature updates that include signatures discovered by WildFire.

Consider purchasing the WildFire subscription service for these additional benefits:

- Sub-hourly (as often as every 15 minutes) WildFire signature updates
- Advanced file type forwarding (APK, Flash, PDF, Microsoft Office, and Java Applet)
- Ability to upload files using the WildFire API
- Ability to forward files to a private WF-500 WildFire appliance

  While the ability to configure a file blocking profile to forward Portable Executable (PE) files to the WildFire cloud for analysis is free, in order to forward files to a private WildFire appliance, a WildFire subscription is required.

| **Enable WildFire** | | |
|---|---|---|
| Step 1 | Confirm that your device is registered and that you have a valid support account as well as any subscriptions you require. | 1. Go to the Palo Alto Networks Support Site, log in, and select **My Devices**.<br>2. Verify that the firewall is listed. If it is not listed, see Register the Firewall.<br>3. (Optional) Activate Licenses and Subscriptions. |

| Enable WildFire (Continued) | | |
|---|---|---|
| Step 2 | Set the WildFire forwarding options.<br><br> If you do not have a WildFire subscription you can only forward executables. | 1. Select **Device > Setup > WildFire** and edit the General Settings.<br>2. (Optional) Specify the **WildFire Server** to which to forward files. By default, the firewall will forward files to the public WildFire cloud hosted in the United States. To forward files to a different WildFire cloud, enter a new value as follows:<br>• To forward to a private WildFire cloud, enter the IP address or FQDN of your WF-500 WildFire appliance.<br>• To forward files to the public WildFire cloud running in Japan, enter `wildfire.paloaltonetworks.jp`.<br>3. (Optional) If you want to change the maximum file size that the firewall can forward for a specific type of file, modify the value in the corresponding field.<br>4. Click **OK** to save your changes. |
| Step 3 | Set up a file blocking profile to forward files to WildFire. | 1. Select **Objects > Security Profiles > File Blocking** and click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the profile.<br>3. Click **Add** to create a forwarding rule and enter a name.<br>4. In the **Action** column, select **forward**.<br>5. Leave the other fields set to **any** to forward any supported file type from any application.<br>6. Click **OK** to save the profile. |
| Step 4 | Attach the file blocking profile to the security policies that allow access to the Internet. | 1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Create Security Rules.<br>2. Click the **Actions** tab within the security policy.<br>3. In the Profile Settings section, click the drop-down and select the file blocking profile you created for WildFire forwarding. (If you don't see a drop-down for selecting a profile, select **Profiles** from the **Profile Type** drop-down. |
| Step 5 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

| Enable WildFire (Continued) | |
|---|---|
| Step 6 Verify that the firewall is forwarding files to WildFire. | 1. Select **Monitor > Logs > Data Filtering**.<br>2. Check the **Action** column for the following actions:<br>• **Forward**— Indicates that the file was successfully forwarded by the file blocking profile attached to the security policy.<br>• **Wildfire-upload-success**—Indicates that the file was sent to WildFire. This means the file is not signed by a trusted file signer and it has not been previously analyzed by WildFire.<br>• **Wildfire-upload-skip**—Indicates that the file was identified as eligible to be sent to WildFire by a file blocking profile/security policy, but did not need to be analyzed by WildFire because it has already been analyzed previously. In this case, the action will display as **forward** in the Data Filtering log because it was a valid forward action, but it was not sent to WildFire and analyzed because the file has already been sent to the WildFire cloud from another session, possibly from another firewall.<br>3. View the WildFire logs by selecting **Monitor > Logs > WildFire Submissions**. If new WildFire logs appear, the firewall is successfully forwarding files to WildFire and WildFire is returning file analysis reports. |

## Scan Traffic for Threats

Security Profiles provide threat protection in security policies. For example, you can apply an antivirus profile to a security policy and all traffic that matches the security policy will be scanned for viruses.

The following sections provide steps for setting up a basic threat prevention configuration:

▲ Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

▲ Set Up File Blocking

### Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

Every Palo Alto Networks next-generation firewall comes with predefined Antivirus, Anti-Spyware, and Vulnerability Protection profiles that you can attach to security policies. There is one predefined Antivirus profile, **default**, which uses the default action for each protocol (block HTTP, FTP, and SMB traffic and alert on SMTP, IMAP, and POP3 traffic). There are two predefined Anti-spyware and Zone Protection profiles:

- **default**—Applies the default action to all client and server critical, high, and medium severity spyware/vulnerability protection events. It does not detect low and informational events.
- **strict**—Applies the block response to all client and server critical, high and medium severity spyware/vulnerability protection events and uses the default action for low and informational events.

To ensure that the traffic entering your network is free from threats, attach the predefined profiles to your basic web access policies. As you monitor the traffic on your network and expand your policy rulebase, you can then design more granular profiles to address your specific security needs.

| Set up Antivirus/Anti-Spyware/Vulnerability Protection | | |
|---|---|---|
| Step 1 | Verify that you have a Threat Prevention license. | • The Threat Prevention license bundles the Antivirus, Anti-Spyware, and the Vulnerability Protection features in one license.<br>• Select **Device > Licenses** to verify that the **Threat Prevention** license is installed and valid (check the expiration date). |
| Step 2 | Download the latest antivirus threat signatures. | 1. Select **Device > Dynamic Updates** and click **Check Now** at the bottom of the page to retrieve the latest signatures.<br>2. In the **Actions** column, click **Download** to install the latest Antivirus, and Applications and Threats signatures. |
| Step 3 | Schedule signature updates.<br> Perform a **download-and-install** on a daily basis for antivirus updates and weekly for applications and threats updates. | 1. From **Device > Dynamic Updates**, click the text to the right of **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.<br>2. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select **Download Only**, you would need to manually go in and click the **Install** link in the **Action** column to install the signature. When you click **OK**, the update is scheduled. No commit is required.<br>3. (Optional) You can also enter the number of hours in the **Threshold** field to indicate the minimum age of a signature before a download will occur. For example, if you entered **10**, the signature must be at least 10 hours old before it will be downloaded, regardless of the schedule.<br>4. In an HA configuration, you can also click the **Sync To Peer** option to synchronize the content update with the HA peer after download/install. This will not push the schedule settings to the peer device, you need to configure the schedule on each device. |

Recommendations for HA Configurations:

- **Active/Passive HA**—If the MGT port is used for antivirus signature downloads, you should configure a schedule on both devices and both devices will download/install independently. If you are using a data port for downloads, the passive device will not perform downloads while it is in the passive state. In this case you would set a schedule on both devices and then select the **Sync To Peer** option. This will ensure that whichever device is active, the updates will occur and will then push to the passive device.
- **Active/Active HA**—If the MGT port is used for antivirus signature downloads on both devices, then schedule the download/install on both devices, but do not select the **Sync To Peer** option. If you are using a data port, schedule the signature downloads on both devices and select **Sync To Peer**. This will ensure that if one device in the active/active configuration goes into the active-secondary state, the active device will download/install the signature and will then push it to the active-secondary device.

 © Palo Alto Networks, Inc.

| Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued) | |
|---|---|
| Step 4 Attach the security profiles to a security policy.<br><br>Attach a clone of a predefined security profile to your basic security policies. That way, if you want to customize the profile you can do so without deleting the read-only predefined **strict** or **default** profile and attaching a customized profile. | 1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab.<br>2. In **Profile Settings**, click the drop-down next to each security profile you would like to enable. In this example we choose default for **Antivirus, Vulnerability Protection,** and **Anti-Spyware**.<br><br>If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.<br><br> 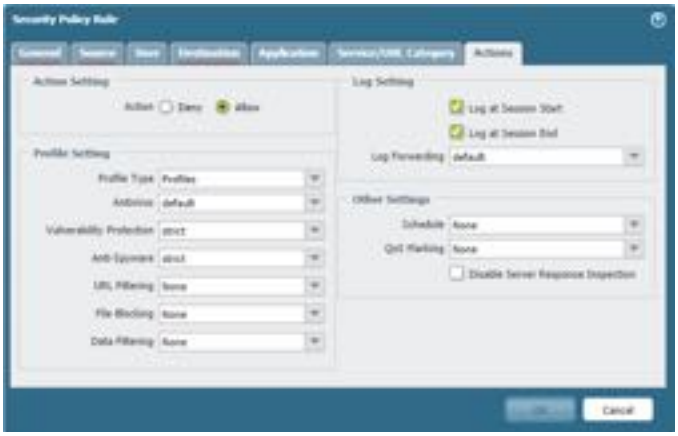 |
| Step 5 Save the configuration. | Click **Commit**. |

## Set Up File Blocking

File Blocking Profiles allow you to identify specific file types that you want to want to block or monitor. The following workflow shows how to set up a basic file blocking profile that prevents users from downloading executable files from the Internet.

| Configure File Blocking | | |
|---|---|---|
| Step 1 | Create the file blocking profile. | 1. Select **Objects > Security Profiles > File Blocking** and click **Add**.<br>2. Enter a **Name** for the file blocking profile, for example `Block_EXE`.<br>3. Optionally enter a **Description**, such as `Block users from downloading exe files from websites`. |
| Step 2 | Configure the file blocking options.<br><br>The **forward** and **continue-and-forward** actions are for forwarding files to WildFire only. | 1. Click **Add** to define the profile settings.<br>2. Enter a **Name**, such as `BlockEXE`.<br>3. Set the **Applications** to which to apply file blocking, or leave it set to **any**.<br>4. Set **File Types** to block. For example, to block download of executables, you would select **exe**.<br>5. Specify the **Direction** in which to block files: **download**, **upload**, or **both**.<br>6. Set the **Action** to one of the following:<br>• **continue**—(web traffic only) Files matching the selected criteria will trigger a customizable response page that requires users to click **Continue** in order to proceed with the download/upload. You must enable response pages on the associated interfaces if you plan to use this option (Step 4).<br>• **block**—Files matching the selected criteria will be blocked from download/upload.<br>• **alert**—Files matching the selected criteria will be allowed, but will generate a log entry in the data filtering log.<br> <br>7. Click **OK** to save the profile. |
| Step 3 | Attach the file blocking profile to the security policies that allow access to content. | 1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Create Security Rules.<br>2. Click the **Actions** tab within the security policy.<br>3. In the Profile Settings section, click the drop-down and select the file blocking profile you created.<br>If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down. |

 © Palo Alto Networks, Inc.

| Configure File Blocking (Continued) | | |
|---|---|---|
| Step 4 | Enable response pages in the management profile for each interface on which you are attaching file blocking profile with a **continue** action. | 1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.<br>2. Select **Response Pages**, as well as any other management services required on the interface.<br>3. Click **OK** to save the interface management profile.<br>4. Select **Network > Interfaces** and select the interface to which to attach the profile.<br>5. On the **Advanced > Other Info** tab, select the interface management profile you just created.<br>6. Click **OK** to save the interface settings. |
| Step 5 | Test the file blocking configuration. | Access a client PC in the trust zone of the firewall and attempt to download an.exe file from a website in the untrust zone. Make sure the file is blocked as expected based on the action you defined in the file blocking profile:<br>• If you selected **alert** as the action, check the data filtering log to make sure you see a log entry for the request.<br>• If you selected **block** as the action, the File Blocking Block Page response page should display.<br>• If you selected the **continue** action, the File Blocking Continue Page response page should display. Click **Continue** to download the file. The following shows the default File Blocking Continue Page.<br> |

## Control Access to Web Content

URL Filtering provides visibility and control over web traffic on your network. With URL filtering enabled, the firewall can categorize web traffic into one or more (from approximately 60) categories. You can then create policies that specify whether to allow, block, or log (alert) traffic based on the category to which it belongs. The following workflow shows how to enable PAN-DB for URL filtering, create security profiles, and attach them to security policies to enforce a basic URL filtering policy.

| Configure URL Filtering | | |
|---|---|---|
| Step 1 | Confirm license information for URL Filtering. | 1. Obtain and install a URL Filtering license. See Activate Licenses and Subscriptions for details.<br>2. Select **Device > Licenses** and verify that the URL Filtering license is valid.<br> |
| Step 2 | Download the seed database and activate the license. | 1. To download the seed database, click **Download** next to **Download Status** in the PAN-DB URL Filtering section of the Licenses page.<br>2. Choose a region (North America, Europe, APAC, Japan) and then click **OK** to start the download.<br>3. After the download completes, click **Activate**.<br> |
| Step 3 | Create a URL filtering profile.<br> Because the default URL filtering profile blocks risky and threat-prone content, clone this profile when creating a new profile in order to preserve the default settings. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select the default profile and then click **Clone**. The new profile will be named default-1.<br>3. Select the new profile and rename it. |

 © Palo Alto Networks, Inc.

| Configure URL Filtering (Continued) | |
|---|---|
| Step 4 Define how to control access to web content.<br><br>If you are not sure what traffic you want to control, consider setting the categories (except for those blocked by default) to alert. You can then use the visibility tools on the firewall, such as the ACC and App Scope, to determine which web categories to restrict to specific groups or to block entirely. You can then go back and modify the profile to block and allow categories as desired.<br><br>You can also define specific sites to always allow or always block regardless of category and enable the safe search option to filter search results when defining the URL Filtering profile. | 1. For each category that you want visibility into or control over, select a value from the **Action** column as follows:<br>• If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **allow**.<br>• For visibility into traffic to sites in a category, select **alert**.<br>• To present a response page to users attempting to access a particular category to alert them to the fact that the content they are accessing might not be work appropriate, select **continue**.<br>• To prevent access to traffic that matches the associated policy, select **block** (this also generates a log entry).<br> <br>2. Click **OK** to save the URL filtering profile. |
| Step 5 Attach the URL filtering profile to a security policy. | 1. Select **Policies > Security**.<br>2. Select the desired policy to modify it and then click the **Actions** tab.<br>3. If this is the first time you are defining a security profile, select **Profiles** from the **Profile Type** drop-down.<br>4. In the **Profile Settings** list, select the profile you just created from the **URL Filtering** drop-down. (If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.)<br>5. Click **OK** to save the profile.<br>6. **Commit** the configuration. |

| Configure URL Filtering (Continued) | | |
|---|---|---|
| Step 6 | Enable Response Pages in the management profile for each interface on which you are filtering web traffic. | 1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.<br>2. Select **Response Pages**, as well as any other management services required on the interface.<br>3. Click **OK** to save the interface management profile.<br>4. Select **Network > Interfaces** and select the interface to which to attach the profile.<br>5. On the **Advanced > Other Info** tab, select the interface management profile you just created.<br>6. Click **OK** to save the interface settings. |
| Step 7 | Save the configuration. | Click **Commit**. |
| Step 8 | Test the URL filtering configuration. | Access a client PC in the trust zone of the firewall and attempt to access a site in a blocked category. Make sure URL filtering is applied based on the action you defined in the URL filtering profile:<br>• If you selected **alert** as the action, check the data filtering log to make sure you see a log entry for the request.<br>• If you selected the **continue** action, the URL Filtering Continue and Override Page response page should display. **Continue** to the site.<br>• If you selected **block** as the action, the URL Filtering and Category Match Block Page response page should display as follows:<br> |

## For More Information

For more detailed information on how to protect your enterprise from threats, see Threat Prevention. For details on how to scan encrypted (SSH or SSL) traffic for threats, see Decryption.

For information about the threats and applications that Palo Alto Networks products can identify, visit the following links:

- Applipedia—Provides details on the applications that Palo Alto Networks can identify.
- Threat Vault—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

 © Palo Alto Networks, Inc.

# Best Practices for Completing the Firewall Deployment

Now that you have integrated the firewall into your network and enabled the basic security features, you can begin configuring more advanced features. Here are some things to consider next:

- Learn about the different Management Interfaces that are available to you and how to access and use them.
- Set up High Availability—High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in a two-device cluster provides redundancy and allows you to ensure business continuity.
- Configure the Master Key—Every Palo Alto Networks firewall has a default master key that encrypts private keys that are used to authenticate administrators when they access management interfaces on the firewall. As a best practice to safeguard the keys, configure the master key on each firewall to be unique.
- Manage Firewall Administrators—Every Palo Alto Networks firewall and appliance is preconfigured with a default administrative account (admin) that provides full read-write access (also known as superuser access) to the device. As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This allows you to better protect the device from unauthorized configuration (or modification) and to enable logging of the actions of each individual administrator.
- Enable User Identification (User-ID)—User-ID is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses.
- Enable Decryption—Palo Alto Networks firewalls provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted or tunneled traffic.
- Enable Passive DNS Collection for Improved Threat Intelligence—Enable this opt-in feature to enable the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis in order to improve threat intelligence and threat prevention capabilities.
- Follow the Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions.

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

# Device Management

Administrators can configure, manage, and monitor the Palo Alto Networks firewalls using the web interface, the CLI, and the API management interfaces. Role-based administrative access to the management interfaces can be customized in order to delegate specific tasks or permissions to certain administrators. See the following topics for information on device management options, including how to begin using the management interfaces and how to customize administrator roles:

- Management Interfaces
- Manage Firewall Administrators
- Reference: Web Interface Administrator Access
- Reference: Port Numbers Used by Palo Alto Networks Devices
- Reset the Firewall to Factory Default Settings

# Management Interfaces

PAN-OS firewalls and Panorama provide three user interfaces: a web interface, a command line interface (CLI), and a XML-based management API. See the following topics for how to access and begin using each of the device management interfaces:

- Use the Web Interface to complete administrative tasks and generate reports from the web interface with relative ease. This graphical interface allows you to access the firewall using HTTPS and it is the best way to perform administrative tasks.
- Use the Command Line Interface (CLI) to type through the commands in rapid succession to complete a series of tasks. The CLI is a no-frills interface that supports two command modes and each mode has its own hierarchy of commands and statements. When you get familiar with the nesting structure and the syntax of the commands, the CLI allows quick response times and offers administrative efficiency.
- Use the XML API to streamline your operations and integrate with existing, internally developed applications and repositories. The XML API is provided as a web service that is implemented using HTTP/HTTPS requests and responses.

 © Palo Alto Networks, Inc.

# Use the Web Interface

The following topics describes how to begin using the firewall web interface. For detailed information about the tabs and fields that are available in the web interface, refer to the Web Interface Reference Guide.


- Launch the Web Interface
- Navigate the Web Interface
- Commit Changes
- Use Configuration Pages
- Required Fields
- Lock Transactions


## Launch the Web Interface

The following web browsers are supported for access to the web interface for PAN-OS firewalls and Panorama:

- Internet Explorer 7+
- Firefox 3.6+
- Safari 5+
- Chrome 11+

Launch an Internet browser and enter the firewall's IP address. Enter your user credentials. If logging in to the firewall for the first time, type the default **admin** into both the **Name** and **Password** fields.

To view information on how to use a specific page and an explanation of the fields and options on the page, click the **Help** icon

in the upper right area of the page to open the online help system. In addition to displaying context-sensitive help for a page, clicking the **Help** icon displays a help navigation pane with options to browse and search all help content.

## Navigate the Web Interface

The following conventions apply when using the web interface.

- To display the menu items for a general functional category, click the tab, such as **Objects** or **Device**, near the top of the browser window.

- Click an item on the side menu to display a panel.

- To display submenu items, click the ⊞ icon to the left of an item. To hide submenu items, click the ⊟ icon to the left of the item.

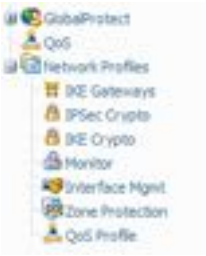


- On most configuration pages, you can click **Add** to create a new item.

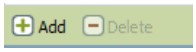


- To delete one or more items, select their check boxes and click **Delete**. In most cases, the system prompts you to confirm by clicking **OK** or to cancel the deletion by clicking **Cancel**.

- On some configuration pages, you can select the check box for an item and click **Clone** to create a new item with the same information as the selected item.

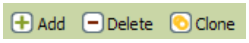


- To modify an item, click its underlined link.

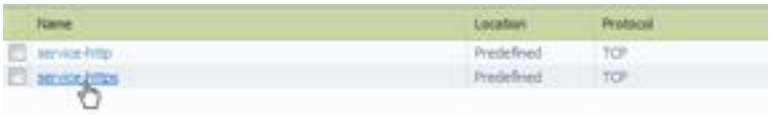


- To view the current list of tasks, click the **Tasks** icon in the lower right corner of the page. The Task Manager window opens to show the list of tasks, along with status, start times, associated messages, and actions. Use the **Show** drop-down to filter the list of tasks.

 © Palo Alto Networks, Inc.

- The web interface language is controlled by the current language of the computer that is managing the device if a specific language preference has not been defined. For example, if the computer you use to manage the firewall has a locale of Spanish, when you log in to the firewall, the web interface will be in Spanish.
- To specify a language that will always be used for a given account regardless of the locale of the computer, click the **Language** icon in the lower right corner of the page and the Language Preference window opens. Click the drop-down to select the desired language and then click **OK** to save your change.

- On pages that list information you can modify (for example, the **Setup** page on the **Devices** tab), click the icon in the upper right corner of a section to edit the settings.

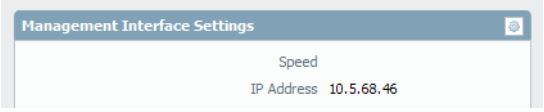


- After you configure settings, you must click **OK** or **Save** to store the changes. When you click **OK**, the current *candidate* configuration is updated.

## Commit Changes

Click **Commit** at the top of the web interface to open the Commit dialog box, which displays the following options. Certain options only display if you click **Advanced**.

If dependencies between the configuration changes you included and excluded cause a validation error, perform the commit with all the changes included. For example, if your changes introduce a new Log Forwarding profile (an object) that references a new Syslog server profile (a device setting), the commit must include both policy and object configurations and device and network configurations.

- **Include Device and Network configuration**—Include the device and network configuration changes in the commit operation.
- **Include Shared Object configuration**—(Multi-virtual system firewalls only) Include the shared object configuration changes in the commit operation.
- **Include Policy and Objects**—(Non-multi-virtual system firewalls only) Include the policy and object configuration changes in the commit operation.
- **Include virtual system configuration**—Include all virtual systems or choose **Select one or more virtual systems**.

- **Preview Changes**—Click this button to bring up a two-pane window that shows proposed changes in the candidate configuration compared to the current running configuration. You can choose the number of lines of context to display, or show all lines. Changes are color coded based on items that you and other administrators added (green), modified (yellow), or deleted (red) since the last commit.

Because the preview results display in a new window, your browser must allow pop-ups. If the preview window does not open, refer to your browser documentation for the steps to unblock pop-ups.

## Use Configuration Pages

The tables on configuration pages include sorting and column chooser options. Click a column header to sort on that column, and click again to change the sort order. Click the arrow to the right of any column and select check boxes to choose the columns to display.

## Required Fields

Required fields are shown with a light yellow background. A message indicating that the field is required appears when you hover over or click in the field entry area.

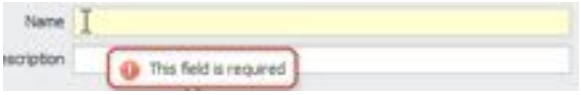


## Lock Transactions

The web interface provides support for multiple administrators by allowing an administrator to lock a current set of transactions, thereby preventing configuration changes or commit operations by another administrator until the lock is removed. The following types of locks are supported:

- **Config lock**—Blocks other administrators from making changes to the configuration. This type of lock can be set globally or for a virtual system. It can be removed only by the administrator who set it or by a superuser on the system.

 © Palo Alto Networks, Inc.

- **Commit Lock**—Blocks other administrators from committing changes until all of the locks have been released. This type of lock prevents collisions that can occur when two administrators are making changes at the same time and the first administrator finishes and commits changes before the second administrator has finished. The lock is released when the current changes are committed by the administrator who applied the lock, or it can be released manually.

Any administrator can open the lock window to view the current transactions that are locked, along with a time stamp for each.

To lock a transaction, click the unlocked icon on the top bar to open the Locks dialog box. Click **Take a Lock**, select the scope of the lock from the drop-down, and click **OK**. Add additional locks as needed, and then click **Close** to close the Lock dialog box.

The transaction is locked, and the icon on the top bar changes to a locked icon that shows the number of locked items in parentheses.

To unlock a transaction, click the locked icon on the top bar to open the Locks window. Click the icon for the lock that you want to remove, and click **Yes** to confirm. Click **Close** to close the Lock dialog box.

You can arrange to automatically acquire a commit lock by selecting the **Automatically acquire commit lock** check box in the Management area of the **Device Setup** page.

# Use the Command Line Interface (CLI)

The PAN-OS CLI allows you to access Firewall and Panorama devices, view status and configuration information, and modify configurations. Access to the PAN-OS CLI is provided through SSH, Telnet, or direct console access.

The following topics describe how to access and begin using the PAN-OS CLI:

- Access the PAN-OS CLI
- Operational and Configuration Modes

For more information on the CLI, refer to the *PAN-OS Command Line Interface Reference Guide*.

## Access the PAN-OS CLI

You can connect to the CLI on the firewall either directly using a serial cable from your management computer to the Console port, or by opening a Secure Shell (SSH) connection to the firewall MGT port as follows:

| **Access the PAN-OS CLI** | |
|---|---|
| Step 1 | Launch a terminal emulator, such as PuTTY, and open a connection to the firewall using one of the following methods:<br>• **Console Connection**—Connect a serial cable from your computer to the Console port and connect use the terminal emulation software to open a Serial connection to the firewall. The firewall requires that you configure the terminal emulation software with the following settings:<br>  • Data rate: 9600<br>  • Data bits: 8<br>  • Parity: none<br>  • Stop bits: 1<br>  • Flow control: None<br>• **SSH Connection**—Use the terminal emulation software to open an SSH connection (default port 22) to the IP address you assigned to the MGT port during initial configuration. |
| Step 2 | Enter the administrative username. The default is admin. |
| Step 3 | Enter the administrative password. The default is admin. |
| Step 4 | The PAN-OS CLI opens in Operational mode, and the CLI prompt is displayed: |
| Step 5 | `username@hostname>` |

 © Palo Alto Networks, Inc.

## Operational and Configuration Modes

When you log in, the PAN-OS CLI opens in Operational mode. You can move between Operational and Configuration modes at any time. Use Operational mode to view the state of the system, navigate the PAN-OS CLI, and enter configuration mode. Use Configuration mode to view and modify the configuration hierarchy.

- To enter Configuration mode from Operational mode, use the `configure` command:

```
username@hostname> configure
Entering configuration mode
[edit]
username@hostname#
```

- To leave Configuration mode and return to Operational mode, use the `quit` or `exit` command:

```
username@hostname# quit
Exiting configuration mode
username@hostname>
```

- To enter an Operational mode command while in Configuration mode, use the `run` command, for example:

```
username@hostname# run ping host 1.1.1.2
PING 1.1.1.1 (1.1.1.1) 56(84) bytes of data
...
username@hostname#
```

- To direct an Operational mode command to a particular VSYS, specify the target VSYS with the following command:

```
username@hostname# set system setting target-vsys <vsys_name>
```

## Use the XML API

Palo Alto Networks XML API uses standard HTTP requests to send and receive data, allowing access to several types of data on the device so the data can be easily integrated with and used in other systems. Use the XML-based Management API to view a firewall or Panorama's configuration, extract report data in XML format, and execute operational commands. API calls can be made directly from command line utilities such as cURL or wget, or using any scripting or application framework that supports RESTful services. When using the API with command lines tools, both HTTP GET and POST methods are supported.

You must generate an API key in order to use the XML API. The API key authenticates the user to the firewall, application, or Panorama. After you have generated an API key, you can use the key to perform device configuration and operational tasks, retrieve reports and logs, and import and export files. See Generate an API Key for steps to generate an API key.

The following table shows the URL structure for API requests:

| PAN-OS Version | XML API URL Structure |
|---|---|
| Prior to PAN-OS 4.1.0 | http(s)://hostname/esp/restapi.esp?request-parameters-values |
| PAN-OS 4.1.0 and later | http(s)://hostname/api/?request-parameters-values |

URL structure item definitions:

- **hostname**—Device's IP address or Domain name.
- **request-parameters-values**—A series of multiple 'parameter=value' pairs separated by the ampersand character (&). These values can either be keywords or data-values in standard or XML format (response data is always in XML format).

There are APIs for PAN-OS, User-ID, and WildFire products. For more information on how to use the API interface, refer to the PAN-OS XML API Usage Guide. To access the online community for developing scripts, visit: https://live.paloaltonetworks.com/community/devcenter.

### Generate an API Key

In order to use the API to manage a firewall or application, an API key is required to authenticate all API calls. Admin account credentials are used to generate API keys.

As a best practice, create a separate admin account for XML-based administration.

| Generate an API key | |
|---|---|
| Step 1 Create an administrator account. | 1. In the web interface, on the **Device > Administrators** tab, click **Add**.<br>2. Enter a login **Name** for the admin.<br>3. Enter and confirm a **Password** for the admin.<br>4. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

| Generate an API key (Continued) | | |
|---|---|---|
| Step 2 | Request an API key. | Replace the hostname, username and password parameters in the following URL with the appropriate values from your administrator account credentials:<br>http(s)://hostname/api/?type=keygen&user=username&password=password<br>The API key is displayed in an XML block. For example:<br><response status="success"><br><result><br><key>0RgWc42Oi0vDx2WRUIUM6A</key><br></result><br></response> |
| Step 3 | (Optional) Revoke or change an API key.<br>For PAN-OS 4.1.0 and later releases, generating an API key using the same administrator account credentials returns unique API keys every time, and all of the keys are valid.<br>You can choose to revoke and then change an API key associated with an administrator account by changing the password associated with the administrator account. Any API keys that were generated using the previous credentials would no longer be valid. | 1. On the **Device > Adminstrators** tab, open the administrator account associated with the API key.<br>2. Enter and confirm a new **Password** for the administrator account.<br>3. Click **OK** and **Commit**.<br>Any API keys associated with the admin account prior to the password change are revoked upon Commit.<br>4. (Optional) Use the updated administrator account credentials to generate a new API key. See Step 2. |

Example work flow using an API key:

Request an API key by entering the URL with the appropriate values in a web browser:

```
https://10.xx.10.50/esp/restapi.esp?type=keygen&user=admin&password=admin
```

Entering the URL displays an XML block that contains the API key:

```
<response status="success">
  <result>
    <key>0RgWc42Oi0vDx2WRUIUM6A</key>
  </result>
</response>
```

Continue to use the API key to create API requests. For example, to generate a report:

```
https://10.xx.10.50/esp/restapi.esp?type=report&reporttype=dynamic&reportname
=top-app-summary&period=last-hour&topn=5&key=0RgWc42Oi0vDx2WRUIUM6A=
```

# Manage Firewall Administrators

Every Palo Alto Networks firewall and appliance is preconfigured with a default administrative account (admin) that provides full read-write access (also known as superuser access) to the device.

As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This allows you to better protect the device from unauthorized configuration (or modification) and to enable logging of the actions of each individual administrator.

The following topics describe the various ways you can set up administrative accounts and provides procedures for setting up basic administrative access:

- Administrative Roles
- Administrative Authentication
- Create an Administrative Account

## Administrative Roles

The way you configure administrator accounts depends on the security requirements within your organization, whether you have existing authentication services you want to integrate with, and how many different administrative roles you require. A *role* defines the type of access the associated administrator has to the system. There are two types of roles you can assign:

- **Dynamic Roles**—These are built-in roles that provide access to the firewall. When new features are added, the firewall automatically updates the definitions of dynamic roles; you never need to manually update them. The following table lists the access privileges associated with dynamic roles.

| Dynamic Role | Privileges |
|---|---|
| Superuser | Full access to the firewall, including defining new administrator accounts and virtual systems. You must have superuser privileges to create an administrative user with superuser privileges. |
| Superuser (read-only) | Read-only access to the firewall. |
| Virtual system administrator | Full access to a selected virtual system (vsys) on the firewall. |
| Virtual system administrator (read-only) | Read-only access to a selected vsys on the firewall. |
| Device administrator | Full access to all firewall settings except for defining new accounts or virtual systems. |
| Device administrator (read-only) | Read-only access to all firewall settings except password profiles (no access) and administrator accounts (only the logged in account is visible). |

- **Admin Role Profiles**—Custom roles you can configure for more granular access control over the functional areas of the web interface, CLI, and XML API. For example, you can create an Admin Role profile for your operations staff that provides access to the firewall and network configuration areas of the web interface and a separate profile for your security administrators that provides access to security policy definitions, logs, and reports. On a multi-vsys firewall, you can select whether the role defines access for all

 © Palo Alto Networks, Inc.

virtual systems or for a specific vsys. When new features are added to the product, you must update the roles with corresponding access privileges: the firewall does not automatically add new features to custom role definitions. For details on the privileges you can configure for custom administrator roles, see Reference: Web Interface Administrator Access.

## Administrative Authentication

There are four ways you can authenticate administrative users:

- **Local administrator account with local authentication**—Both the administrator account credentials and the authentication mechanisms are local to the firewall. You can further secure the local administrator account by creating a password profile that defines a validity period for passwords and by setting device-wide password complexity settings.
- **Local administrator account with SSL-based authentication**—With this option, you create the administrator accounts on the firewall, but authentication is based on SSH certificates (for CLI access) or client certificates/common access cards (for the web interface). Refer to the article How to Configure Certificate-based Authentication for the WebUI for details on how to configure this type of administrative access.
- **Local administrator account with external authentication**—The administrator accounts are managed on the local firewall, but the authentication functions are offloaded to an existing LDAP, Kerberos, or RADIUS service. To configure this type of account, you must first create an authentication profile that defines how to access the external authentication service and then create an account for each administrator that references the profile.
- **External administrator account and authentication**—Account administration and authentication are handled by an external RADIUS server. To use this option, you must define Vendor Specific Attributes (VSAs) on your RADIUS server that map to the admin role and, optionally, the virtual system objects you have defined on the Palo Alto Networks device. Refer to the Radius Vendor Specific Attributes (VSA) article for details on how to configure this type of administrative access.

## Create an Administrative Account

Create administrative accounts to define access and administrative privileges for firewall administrators. Because it is common to delegate specific administrative tasks to specific administrators with varying roles, Palo Alto Networks recommends that you create admin role profiles that allow administrators access only to the areas of the management interface that are required to perform their jobs. You can assign the various roles you create to individual administrator accounts and specify access privileges to each management interface: the web interface, the Command Line Interface (CLI), and the XML-based Management API. By creating admin roles with very granular access privileges, you can ensure that sensitive company data is protected and end user privacy is ensured.

The following procedure describes how to create a local administrator account with local authentication, including how to set administrator access for each management interface.

© Palo Alto Networks, Inc.

| Create a Local Administrator | |
|---|---|
| Step 1 Create the Admin Role Profiles that you plan to assign to your administrators (this does not apply if you plan to use Dynamic Roles). The Admin Role Profiles define what type of access to give to the different sections of the web interface, CLI, and XML API for each administrator that you assign a role to.<br><br>You can use this step to set particularly granular access privileges for web interface users. For details on what a specific option enables on the **Web UI** tab, see Web Interface Access Privileges. | Complete the following steps for each role you want to create:<br>1. Select **Device > Admin Roles** and then click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the role.<br>3. On the Web UI, Command Line and/or XML API tabs, specify the access to allow for each management interface:<br>• On the **Web UI** and/or **XML API** tabs, set the access levels for each functional area of the interface by clicking the icon to toggle it to the desired setting: Enable, Read Only, or Disable.<br>• On the **Command Line** tab, specify the type of access to allow to the CLI: **superreader**, **deviceadmin**, or **devicereader** (for Device roles); **vsysadmin** or **vsysreader** (for Virtual System roles); or **None** to disable CLI access entirely.<br>4. Click **OK** to save the profile.<br><br>For example, allow an admin full access to a device using the XML API, with the exception of importing or exporting files:<br> |

 © Palo Alto Networks, Inc.

| Create a Local Administrator (Continued) | | |
|---|---|---|
| Step 2 | (Optional) Set requirements for local user-defined passwords. | • **Create Password Profiles**—Define how often administrators must change their passwords. You can create multiple password profiles and apply them to administrator accounts as needed to enforce the desired security. To create a password profile, select **Device > Password Profiles** and then click **Add**.<br>• **Configure minimum password complexity settings**—Define rules that govern password complexity, allowing you to force administrators to create passwords that are harder to guess, crack, or compromise. Unlike password profiles, which can be applied to individual accounts, these rules are device-wide and apply to all passwords. To configure the settings, select **Device > Setup** and then Edit in the Minimum Password Complexity section. |
| Step 3 | Create an account for each administrator. | 1. Select **Device > Administrators** and then click **Add**.<br>2. Enter a user **Name** and **Password** for the administrator, or create an **Authentication Profile** to use for validating an administrative user's credentials to an external authentication server. See Step 4 for details on setting up an authentication profile.<br>3. Select the **Role** to assign to this administrator. You can either select one of the predefined Dynamic roles or a custom Role Based profile if you created one in Step 1.<br>4. (Optional) Select a **Password Profile**.<br>5. Click **OK** to save the account. |

| Create a Local Administrator (Continued) | | |
|---|---|---|
| Step 4 | (Optional) Set up authentication to an external server—LDAP, RADIUS or Kerberos.<br><br>The server profile specifies how the firewall can connect to the authentication service you plan to use. | 1. Select **Device > Authentication Profile** and then click **Add**.<br>2. Enter a user **Name** to identify the authentication profile.<br>3. Define the conditions for locking out the administrative user.<br>a. Enter the **Lockout Time**. This is the number of minutes that a user is locked out upon reaching the maximum number of failed attempts (0-60 minutes; default 0). 0 means that the lockout is in effect until it is manually unlocked.<br>b. Enter the **Failed Attempts** count. This is the number of failed login attempts that are allowed before the account is locked out (1-10; default 0). By default, the failed attempt count is 0 and the user is not locked out despite repeated failure to authenticate.<br>4. Specify the users and groups that are explicitly allowed to authenticate. By adding an **Allow List** to an authentication profile, you can limit access to specific users in a user group/directory.<br>• Select the **All** check box to allow all users.<br>• Click **Add** and enter the first few characters of a name in the field to list all the users and user groups that start with those characters. Repeat to add as many users/user groups as required.<br>5. In the **Authentication** drop-down, select the type of authentication you plan to use on your network.<br>If you plan to use local database authentication, you must create the local database. Select **Device > Local User Database** and add the users and groups to be authenticated.<br>6. For access to an external authentication server (not local database), select the appropriate server profile in the **Server Profile** drop-down. To create a new server profile, click the link next to New and continue with configuring access to the **LDAP**, **RADIUS** or **Kerberos** server.<br>7. Click **OK**. |
| Step 5 | Commit your changes. | 1. Click **Commit**. |


© Palo Alto Networks, Inc.

# Reference: Web Interface Administrator Access

While the privileges associated with dynamic administrator roles are fixed, you can configure privileges at a granular level for custom administrator roles. Configuring privileges at a granular level ensures that lower level administrators cannot access certain information. You can create custom roles for firewall administrators (see Create an Administrative Account), Panorama administrators, or Device Group and Template administrators (refer to the Panorama Administrator's Guide). The following topics describe the privileges you can configure for custom administrator roles.

- Web Interface Access Privileges
- Panorama Web Interface Access

## Web Interface Access Privileges

If you want to prevent a role-based administrator from accessing specific tabs on the web interface, you can disable the tab and the administrator will not even see it when logging in using the associated role-based administrative account. For example, you could create an Admin Role Profile for your operations staff that provides access to the **Device** and **Network** tabs only and a separate profile for your security administrators that provides access to the **Object**, **Policy**, and **Monitor** tabs.

The following table describes the tab-level access privileges you can assign to the admin role profile. It also provides cross-references to additional tables that detail granular privileges within a tab. For specific information on how to set the admin role profile to protect end user privacy, see Define User Privacy Settings in the Admin Role Profile.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Dashboard | Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets. | Yes | No | Yes |
| ACC | Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Yes | No | Yes |
| Monitor | Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the admin can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab. | Yes | No | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Policies | Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the admin can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab. | Yes | No | Yes |
| Objects | Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab. | Yes | No | Yes |
| Network | Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the admin can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab. | Yes | No | Yes |
| Device | Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any device-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab.<br> You cannot enable access to the **Admin Roles** or **Administrators** nodes for a role-based administrator even if you enable full access to the **Device** tab. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

## Provide Granular Access to the Monitor Tab

In some cases you might want to enable the administrator to view some but not all areas of the Monitor tab. For example, you may want to restrict operations admins to the configuration and system logs only, because they do not contain sensitive user data. Although this section of the admin role definition specifies what areas of the Monitor tab the admin can see, you can also couple privileges in this section with privacy privileges, such as disabling the ability to see user names in logs and reports. One thing to keep in mind, however, is that any system-generated reports will still show user names and IP addresses even if you disable that functionality in the role. For this reason, if you do not want the admin to see any of the private user information, you should disable access to the specific reports as detailed in the following table.

The following table lists the **Monitor** tab access levels and the administrator roles for which they are available.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Monitor | Enables or disables access to the **Monitor** tab. If disabled, the admin will not see this tab or any of the associated logs or reports. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Logs | Enables or disables access to all log files. You can also leave this privilege enabled and then disable specific logs that you do not want the admin to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the logs, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Traffic | Specifies whether the admin can see the traffic logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Threat | Specifies whether the admin can see the threat logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| URL Filtering | Specifies whether the admin can see the URL filtering logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| WildFire Submissions | Specifies whether the admin can see the WildFire logs. These logs are only available if you have a WildFire subscription. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Data Filtering | Specifies whether the admin can see the data filtering logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| HIP Match | Specifies whether the admin can see the HIP Match logs. HIP Match logs are only available if you have a GlobalProtect portal license and gateway subscription. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Configuration | Specifies whether the admin can see the configuration logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| System | Specifies whether the admin can see the system logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Alarms | Specifies whether the admin can see system-generated alarms. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Packet Capture | Specifies whether the admin can see packet captures (pcaps) from the **Monitor** tab. Keep in mind that packet captures are raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the pcap and you should therefore disable the Packet Capture privilege if you are concerned about user privacy. | Firewall: Yes<br>Panorama: No<br>Device Group/Template: No | Yes | Yes | Yes |
| App Scope | Specifies whether the admin can see the App Scope visibility and analysis tools. Enabling App Scope enables access to all of the **App Scope** charts. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Session Browser | Specifies whether the admin can browse and filter current running sessions on the firewall. Keep in mind that the session browser shows raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the session browser and you should therefore disable the **Session Browser** privilege if you are concerned about user privacy. | Firewall: Yes<br>Panorama: No<br>Device Group/Template: No | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Botnet | Specifies whether the admin can generate and view botnet analysis reports or view botnet reports in read-only mode. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in scheduled botnet reports and you should therefore disable the **Botnet** privilege if you are concerned about user privacy. | Firewall: Yes<br>Panorama: No<br>Device Group/Template: No | Yes | Yes | Yes |
| PDF Reports | Enables or disables access to all PDF reports. You can also leave this privilege enabled and then disable specific PDF reports that you do not want the admin to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Manage PDF Summary | Specifies whether the admin can view, add or delete PDF summary report definitions. With read-only access, the admin can see PDF summary report definitions, but not add or delete them. If you disable this option, the admin can neither view the report definitions nor add/delete them. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| PDF Summary Reports | Specifies whether the admin can see the generated PDF Summary reports in **Monitor > Reports**. If you disable this option, the **PDF Summary Reports** category will not display in the **Reports** node. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| User Activity Report | Specifies whether the admin can view, add or delete User Activity report definitions and download the reports. With read-only access, the admin can see User Activity report definitions, but not add, delete, or download them. If you disable this option, the admin cannot see this category of PDF report. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Report Groups | Specifies whether the admin can view, add or delete report group definitions. With read-only access, the admin can see report group definitions, but not add or delete them. If you disable this option, the admin cannot see this category of PDF report. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Email Scheduler | Specifies whether the admin can schedule report groups for email. Because the generated reports that get emailed may contain sensitive user data that is not removed by disabling the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** options and because they may also show log data to which the admin does not have access, you should disable the **Email Scheduler** option if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Manage Custom Reports | Enables or disables access to all Custom report functionality. You can also leave this privilege enabled and then disable specific custom report categories that you do not want the admin to be able to access. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option.<br> Reports that are scheduled to run rather than run on demand will show IP address and user information. In this case, be sure to restrict access to the corresponding report areas. In addition, the custom report feature does not restrict the ability to generate reports that contain log data contained in logs that are excluded from the admin role. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Application Statistics | Specifies whether the admin can create a custom report that includes data from the application statistics database. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Data Filtering Log | Specifies whether the admin can create a custom report that includes data from the data filtering log. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Threat Log | Specifies whether the admin can create a custom report that includes data from the threat log. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Threat Summary | Specifies whether the admin can create a custom report that includes data from the threat summary database. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Traffic Log | Specifies whether the admin can create a custom report that includes data from the traffic log. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Traffic Summary | Specifies whether the admin can create a custom report that includes data from the traffic summary database. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Url Log | Specifies whether the admin can create a custom report that includes data from the URL filtering log. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Hipmatch | Specifies whether the admin can create a custom report that includes data from the HIP match log. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| View Scheduled Custom Reports | Specifies whether the admin can view a custom report that has been scheduled to generate. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| View Predefined Application Reports | Specifies whether the admin can view Application Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| View Predefined Threat Reports | Specifies whether the admin can view Threat Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| View Predefined URL Filtering Reports | Specifies whether the admin can view URL Filtering Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| View Predefined Traffic Reports | Specifies whether the admin can view Traffic Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |

## Provide Granular Access to the Policy Tab

If you enable the Policy option in the admin role profile, you can then enable, disable, or provide read-only access to specific nodes within the tab as necessary for the admin role you are defining. By enabling access to a specific policy type, you enable the ability to view, add, or delete policy rules. By enabling read-only access to a specific policy, you enable the admin to view the corresponding policy rule base, but not add or delete rules. Disabling access to a specific type of policy prevents the admin from seeing the policy rule base.

Because policy that is based on specific users (by user name or IP address) must be explicitly defined, privacy settings that disable the ability to see full IP addresses or user names do not apply to the Policy tab. Therefore, you should only allow access to the Policy tab to administrators that are excluded from user privacy restrictions.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Security | Enable this privilege to allow the admin to view, add, and/or delete security policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the security policy rule base, disable this privilege. | Yes | Yes | Yes |
| NAT | Enable this privilege to allow the admin to view, add, and/or delete NAT policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the NAT policy rule base, disable this privilege. | Yes | Yes | Yes |
| QoS | Enable this privilege to allow the admin to view, add, and/or delete QoS policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the QoS policy rule base, disable this privilege. | Yes | Yes | Yes |
| Policy Based Forwarding | Enable this privilege to allow the admin to view, add, and/or delete Policy-Based Forwarding (PBF) policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the PBF policy rule base, disable this privilege. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Decryption | Enable this privilege to allow the admin to view, add, and/or delete decryption policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the decryption policy rule base, disable this privilege. | Yes | Yes | Yes |
| Application Override | Enable this privilege to allow the admin to view, add, and/or delete application override policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the application override policy rule base, disable this privilege. | Yes | Yes | Yes |
| Captive Portal | Enable this privilege to allow the admin to view, add, and/or delete captive portal policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the captive portal policy rule base, disable this privilege. | Yes | Yes | Yes |
| DoS Protection | Enable this privilege to allow the admin to view, add, and/or delete DoS protection policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the DoS protection policy rule base, disable this privilege. | Yes | Yes | Yes |

## Provide Granular Access to the Objects Tab

An *object* is a container that groups specific policy filter values—such as IP addresses, URLs, applications, or services—for simplified rule definition. For example, an address object might contain specific IP address definitions for the web and application servers in your DMZ zone.

When deciding whether to allow access to the objects tab as a whole, determine whether the admin will have policy definition responsibilities. If not, the admin probably does not need access to the tab. If, however, the admin will need to create policy, you can enable access to the tab and then provide granular access privileges at the node level.

By enabling access to a specific node, you give the admin the privilege to view, add, and delete the corresponding object type. Giving read-only access allows the admin to view the already defined objects, but not create or delete any. Disabling a node prevents the admin from seeing the node in the web interface.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Addresses | Specifies whether the admin can view, add, or delete address objects for use in security policy. | Yes | Yes | Yes |
| Address Groups | Specifies whether the admin can view, add, or delete address group objects for use in security policy. | Yes | Yes | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Regions | Specifies whether the admin can view, add, or delete regions objects for use in security, decryption, or DoS policy. | Yes | Yes | Yes |
| Applications | Specifies whether the admin can view, add, or delete application objects for use in policy. | Yes | Yes | Yes |
| Application Groups | Specifies whether the admin can view, add, or delete application group objects for use in policy. | Yes | Yes | Yes |
| Application Filters | Specifies whether the admin can view, add, or delete application filters for simplification of repeated searches. | Yes | Yes | Yes |
| Services | Specifies whether the admin can view, add, or delete service objects for use in creating policies that limit the port numbers an application can use. | Yes | Yes | Yes |
| Service Groups | Specifies whether the admin can view, add, or delete service group objects for use in security policy. | Yes | Yes | Yes |
| Tags (Panorama only) | Specifies whether the admin can view, add, or delete tags that have been defined on the device. | Yes | Yes | Yes |
| GlobalProtect | Specifies whether the admin can view, add, or delete HIP objects and profiles. You can restrict access to both types of objects at the GlobalProtect level, or provide more granular control by enabling the GlobalProtect privilege and restricting HIP Object or HIP Profile access. | Yes | No | Yes |
| HIP Objects | Specifies whether the admin can view, add, or delete HIP objects, which are used to define HIP profiles. HIP Objects also generate HIP Match logs. | Yes | Yes | Yes |
| HIP Profiles | Specifies whether the admin can view, add, or delete HIP Profiles for use in security policy and/or for generating HIP Match logs. | Yes | Yes | Yes |
| Dynamic Block Lists | Specifies whether the admin can view, add, or delete dynamic block lists for use in security policy. | Yes | Yes | Yes |
| Custom Objects | Specifies whether the admin can see the custom spyware and vulnerability signatures. You can restrict access to either enable or disable access to all custom signatures at this level, or provide more granular control by enabling the Custom Objects privilege and then restricting access to each type of signature. | Yes | No | Yes |
| Data Patterns | Specifies whether the admin can view, add, or delete custom data pattern signatures for use in creating custom vulnerability protection profiles. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Spyware | Specifies whether the admin can view, add, or delete custom spyware signatures for use in creating custom vulnerability protection profiles. | Yes | Yes | Yes |
| Vulnerability | Specifies whether the admin can view, add, or delete custom vulnerability signatures for use in creating custom vulnerability protection profiles. | Yes | Yes | Yes |
| URL Category | Specifies whether the admin can view, add, or delete custom URL categories for use in policy. | Yes | Yes | Yes |
| Security Profiles | Specifies whether the admin can see security profiles. You can restrict access to either enable or disable access to all security profiles at this level, or provide more granular control by enabling the Security Profiles privilege and then restricting access to each type of profile. | Yes | No | Yes |
| Antivirus | Specifies whether the admin can view, add, or delete antivirus profiles. | Yes | Yes | Yes |
| Anti-Spyware | Specifies whether the admin can view, add, or delete anti-spyware profiles. | Yes | Yes | Yes |
| Vulnerability Protection | Specifies whether the admin can view, add, or delete vulnerability protection profiles. | Yes | Yes | Yes |
| URL Filtering | Specifies whether the admin can view, add, or delete URL filtering profiles. | Yes | Yes | Yes |
| File Blocking | Specifies whether the admin can view, add, or delete file blocking profiles. | Yes | Yes | Yes |
| Data Filtering | Specifies whether the admin can view, add, or delete data filtering profiles. | Yes | Yes | Yes |
| DoS Protection | Specifies whether the admin can view, add, or delete DoS protection profiles. | Yes | Yes | Yes |
| Security Profile Groups | Specifies whether the admin can view, add, or delete security profile groups. | Yes | Yes | Yes |
| Log Forwarding | Specifies whether the admin can view, add, or delete log forwarding profiles. | Yes | Yes | Yes |
| Decryption Profile | Specifies whether the admin can view, add, or delete decryption profiles. | Yes | Yes | Yes |
| Schedules | Specifies whether the admin can view, add, or delete schedules for limiting a security policy to a specific date and/or time range. | Yes | Yes | Yes |

© Palo Alto Networks, Inc.

## Provide Granular Access to the Network Tab

When deciding whether to allow access to the **Network** tab as a whole, determine whether the admin will have network administration responsibilities, including GlobalProtect administration. If not, the admin probably does not need access to the tab.

You can also define access to the **Network** tab at the node level. By enabling access to a specific node, you give the admin the privilege to view, add, and delete the corresponding network configurations. Giving read-only access allows the admin to view the already-defined configuration, but not create or delete any. Disabling a node prevents the admin from seeing the node in the web interface.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Interfaces | Specifies whether the admin can view, add, or delete interface configurations. | Yes | Yes | Yes |
| Zones | Specifies whether the admin can view, add, or delete zones. | Yes | Yes | Yes |
| VLANs | Specifies whether the admin can view, add, or delete VLANs. | Yes | Yes | Yes |
| Virtual Wires | Specifies whether the admin can view, add, or delete virtual wires. | Yes | Yes | Yes |
| Virtual Routers | Specifies whether the admin can view, add, modify or delete virtual routers. | Yes | Yes | Yes |
| IPSec Tunnels | Specifies whether the admin can view, add, modify, or delete IPSec Tunnel configurations. | Yes | Yes | Yes |
| DHCP | Specifies whether the admin can view, add, modify, or delete DHCP server and DHCP relay configurations. | Yes | Yes | Yes |
| DNS Proxy | Specifies whether the admin can view, add, modify, or delete DNS proxy configurations. | Yes | Yes | Yes |
| GlobalProtect | Specifies whether the admin can view, add, modify GlobalProtect portal and gateway configurations. You can disable access to the GlobalProtect functions entirely, or you can enable the GlobalProtect privilege and then restrict the role to either the portal or gateway configuration areas. | Yes | No | Yes |
| Portals | Specifies whether the admin can view, add, modify, or delete GlobalProtect portal configurations. | Yes | Yes | Yes |
| Gateways | Specifies whether the admin can view, add, modify, or delete GlobalProtect gateway configurations. | Yes | Yes | Yes |
| MDM | Specifies whether the admin can view add, modify, or delete GlobalProtect MDM server configurations. | Yes | Yes | Yes |
| QoS | | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Network Profiles | Sets the default state to enable or disable for all of the Network settings described below. | Yes | No | Yes |
| IKE Gateways | Controls access to the **Network Profiles > IKE Gateways** node. If you disable this privilege, the administrator will not see the **IKE Gateways** node or define gateways that include the configuration information necessary to perform IKE protocol negotiation with peer gateway.<br>If the privilege state is set to read only, you can view the currently configured IKE Gateways but cannot add or edit gateways. | Yes | Yes | Yes |
| IPSec Crypto | Controls access to the **Network Profiles > IPSec Crypto** node. If you disable this privilege, the administrator will not see the **Network Profiles > IPSec Crypto** node or specify protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPSec SA negotiation.<br>If the privilege state is set to read only, you can view the currently configured IPSec Crypto configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| IKE Crypto | Controls how devices exchange information to ensure secure communication. Specify the protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPsec SA negotiation (IKEv1 Phase-1). | Yes | Yes | Yes |
| Monitor | Controls access to the **Network Profiles > Monitor** node. If you disable this privilege, the administrator will not see the **Network Profiles > Monitor** node or be able to create or edit a monitor profile that is used to monitor IPSec tunnels and monitor a next-hop device for policy-based forwarding (PBF) rules.<br>If the privilege state is set to read only, you can view the currently configured monitor profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| Interface Mgmt | Controls access to the **Network Profiles > Interface Mgmt** node. If you disable this privilege, the administrator will not see the **Network Profiles > Interface Mgmt** node or be able to specify the protocols that are used to manage the firewall.<br>If the privilege state is set to read only, you can view the currently configured Interface management profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Zone Protection | Controls access to the **Network Profiles > Zone Protection** node. If you disable this privilege, the administrator will not see the **Network Profiles > Zone Protection** node or be able to configure a profile that determines how the firewall responds to attacks from specified security zones.<br>If the privilege state is set to read only, you can view the currently configured Zone Protection profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| QoS Profile | Controls access to the **Network Profiles > QoS** node. If you disable this privilege, the administrator will not see the **Network Profiles > QoS** node or be able to configure a profile QoS profile that determines how QoS traffic classes are treated.<br>If the privilege state is set to read only, you can view the currently configured QoS profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |

## Provide Granular Access to the Device Tab

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Setup | Controls access to the **Setup** node. If you disable this privilege, the administrator will not see the **Setup** node or have access to device-wide setup configuration information, such as Management, Operations, Service, Content-ID, Wildfire or Session setup information.<br>If the privilege state is set to read only, you can view the current configuration but cannot make any changes. | Yes | Yes | Yes |
| Config Audit | Controls access to the **Config Audit** node. If you disable this privilege, the administrator will not see the **Config Audit** node or have access to any device-wide configuration information. | Yes | No | Yes |
| Admin Roles | Controls access to the **Admin Roles** node. This function can only be allowed for read only access.<br>If you disable this privilege, the administrator will not see the **Admin Roles** node or have access to any device-wide information concerning admin roles configuration.<br>If you set this privilege to read only, you can view the configuration information for all admin roles configured on the device. | No | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Administrators | Controls access to the **Administrators** node. This function can only be allowed for read only access.<br>If you disable this privilege, the administrator will not see the **Administrators** node or have access to information about their own admin account.<br>If you set this privilege to read only, the administrator can view the configuration information for their own admin account. They will not see any information about other admin accounts configured on the device. | No | Yes | Yes |
| Virtual Systems | Controls access to the **Virtual Systems** node. If you disable this privilege, the administrator will not see or be able to configure virtual systems.<br>If the privilege state is set to read only, you can view the currently configured virtual systems but cannot add or edit a configuration. | Yes | Yes | Yes |
| Shared Gateways | Controls access to the **Shared Gateways** node. Shared gateways allow virtual systems to share a common interface for external communications.<br>If you disable this privilege, the administrator will not see or be able to configure shared gateways.<br>If the privilege state is set to read only, you can view the currently configured shared gateways but cannot add or edit a configuration. | Yes | Yes | Yes |
| User Identification | Controls access to the **User Identification** node. If you disable this privilege, the administrator will not see the **User Identification** node or have access to device-wide User Identification configuration information, such as User Mapping, User-ID Agents, Service, Terminal Services Agents, Group Mappings Settings or Captive Portal Settings.<br>If you set this privilege to read only, the administrator can view configuration information for the device but is not allowed to perform any configuration procedures. | Yes | Yes | Yes |
| VM Information Source | Controls access to the **VM Information Source** node that allows you to configure the firewall/Windows User-ID agent to collect VM inventory automatically. If you disable this privilege, the administrator will not see the **VM Information Source** node.<br>If you set this privilege to read only, the administrator can view the VM information sources configured but cannot add, edit, or delete any sources.<br>This privilege is not available to Device Group and Template administrators. | Yes | Yes | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| High Availability | Controls access to the **High Availability** node. If you disable this privilege, the administrator will not see the **High Availability** node or have access to device-wide high availability configuration information such as General setup information or Link and Path Monitoring.<br>If you set this privilege to read only, the administrator can view High Availability configuration information for the device but is not allowed to perform any configuration procedures. | Yes | Yes | Yes |
| Certificate Management | Sets the default state to enable or disable for all of the Certificate settings described below. | Yes | No | Yes |
| Certificates | Controls access to the **Certificates** node. If you disable this privilege, the administrator will not see the **Certificates** node or be able to configure or access information regarding Device Certificates or Default Trusted Certificate Authorities.<br>If you set this privilege to read only, the administrator can view Certificate configuration information for the device but is not allowed to perform any configuration procedures. | Yes | Yes | Yes |
| Certificate Profile | Controls access to the **Certificate Profile** node. If you disable this privilege, the administrator will not see the **Certificate Profile** node or be able to create certificate profiles.<br>If you set this privilege to read only, the administrator can view Certificate Profiles that are currently configured for the device but is not allowed to create or edit a certificate profile. | Yes | Yes | Yes |
| OCSP Responder | Controls access to the **OCSP Responder** node. If you disable this privilege, the administrator will not see the **OCSP Responder** node or be able to define a server that will be used to verify the revocation status of certificates issues by the PAN-OS device.<br>If you set this privilege to read only, the administrator can view the **OCSP Responder** configuration for the device but is not allowed to create or edit an OCSP responder configuration. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Response Pages | Controls access to the **Response Pages** node. If you disable this privilege, the administrator will not see the **Response Page** node or be able to define a custom HTML message that is downloaded and displayed instead of a requested web page or file.<br>If you set this privilege to read only, the administrator can view the **Response Page** configuration for the device but is not allowed to create or edit a response page configuration. | Yes | Yes | Yes |
| Log Settings | Sets the default state to enable or disable for all of the Log settings described below. | Yes | No | Yes |
| System | Controls access to the **Log Settings > System** node. If you disable this privilege, the administrator will not see the **Log Settings > System** node or be able to specify the severity levels of the system log entries that are logged remotely with Panorama and sent as SNMP traps, syslog messages, and/or email notifications.<br>If you set this privilege to read only, the administrator can view the **Log Settings > System** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |
| Config | Controls access to the **Log Settings > Config** node. If you disable this privilege, the administrator will not see the **Log Settings > Config** node or be able to specify the configuration log entries that are logged remotely with Panorama, and sent as syslog messages and/or email notification.<br>If you set this privilege to read only, the administrator can view the **Log Settings > Config** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |
| HIP Match | Controls access to the **Log Settings > HIP Match** node. If you disable this privilege, the administrator will not see the **Log Settings > HIP Match** node or be able to specify the Host Information Profile (HIP) match log settings that are used to provide information on security policies that apply to GlobalProtect clients<br>If you set this privilege to read only, the administrator can view the **Log Settings > HIP** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Alarms | Controls access to the **Log Settings > Alarms** node. If you disable this privilege, the administrator will not see the **Log Settings > Alarms** node or be able to configure notifications that are generated when a security rule (or group of rules) has been hit repeatedly in a set period of time.<br>If you set this privilege to read only, the administrator can view the **Log Settings > Alarms** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |
| Manage Logs | Controls access to the **Log Settings > Manage Logs** node. If you disable this privilege, the administrator will not see the **Log Settings > Manage Logs** node or be able to clear the indicated logs.<br>If you set this privilege to read only, the administrator can view the **Log Settings > Manage Logs** information but cannot clear any of the logs. | Yes | Yes | Yes |
| Server Profiles | Sets the default state to enable or disable for all of the Server Profiles settings described below. | Yes | No | Yes |
| SNMP Trap | Controls access to the **Server Profiles > SNMP Trap** node. If you disable this privilege, the administrator will not see the **Server Profiles > SNMP Trap** node or be able to specify one or more SNMP trap destinations to be used for system log entries.<br>If you set this privilege to read only, the administrator can view the **Server Profiles > SNMP Trap Logs** information but cannot specify SNMP trap destinations. | Yes | Yes | Yes |
| Syslog | Controls access to the **Server Profiles > Syslog** node. If you disable this privilege, the administrator will not see the **Server Profiles > Syslog** node or be able to specify one or more syslog servers.<br>If you set this privilege to read only, the administrator can view the **Server Profiles > Syslog** information but cannot specify syslog servers. | Yes | Yes | Yes |
| Email | Controls access to the **Server Profiles > Email** node. If you disable this privilege, the administrator will not see the **Server Profiles > Email** node or be able to configure an email profile that can be used to enable email notification for system and configuration log entries<br>If you set this privilege to read only, the administrator can view the **Server Profiles > Email** information but cannot configure and email profile. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Netflow | Controls access to the **Server Profiles > Netflow** node. If you disable this privilege, the administrator will not see the **Server Profiles > Netflow** node or be able to define a NetFlow server profile, which specifies the frequency of the export along with the NetFlow servers that will receive the exported data.<br>If you set this privilege to read only, the administrator can view the **Server Profiles > Netflow** information but cannot define a Netflow profile. | Yes | Yes | Yes |
| RADIUS | Controls access to the **Server Profiles > RADIUS** node. If you disable this privilege, the administrator will not see the **Server Profiles > RADIUS** node or be able to configure settings for the RADIUS servers that are identified in authentication profiles.<br>If you set this privilege to read only, the administrator can view the **Server Profiles > RADIUS** information but cannot configure settings for the RADIUS servers. | Yes | Yes | Yes |
| LDAP | Controls access to the **Server Profiles > LDAP** node. If you disable this privilege, the administrator will not see the **Server Profiles > LDAP** node or be able to configure settings for the LDAP servers to use for authentication by way of authentication profiles.<br>If you set this privilege to read only, the administrator can view the **Server Profiles > LDAP** information but cannot configure settings for the LDAP servers. | Yes | Yes | Yes |
| Kerberos | Controls access to the **Server Profiles > Kerberos** node. If you disable this privilege, the administrator will not see the **Server Profiles > Kerberos** node or configure a Kerberos server that allows users to authenticate natively to a domain controller.<br>If you set this privilege to read only, the administrator can view the **Server Profiles > Kerberos** information but cannot configure settings for Kerberos servers. | Yes | Yes | Yes |
| Local User Database | Sets the default state to enable or disable for all of the Local User Database settings described below. | Yes | No | Yes |
| Users | Controls access to the **Local User Database > Users** node. If you disable this privilege, the administrator will not see the **Local User Database > Users** node or set up a local database on the firewall to store authentication information for remote access users, device administrators, and captive portal users.<br>If you set this privilege to read only, the administrator can view the **Local User Database > Users** information but cannot set up a local database on the firewall to store authentication information. | Yes | Yes | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| User Groups | Controls access to the **Local User Database > Users** node. If you disable this privilege, the administrator will not see the **Local User Database > Users** node or be able to add user group information to the local database.<br>If you set this privilege to read only, the administrator can view the **Local User Database > Users** information but cannot add user group information to the local database. | Yes | Yes | Yes |
| Authentication Profile | Controls access to the **Authentication Profile** node. If you disable this privilege, the administrator will not see the **Authentication Profile** node or be able to create or edit authentication profiles that specify local database, RADIUS, LDAP, or Kerberos settings that can be assigned to administrator accounts.<br>If you set this privilege to read only, the administrator can view the **Authentication Profile** information but cannot create or edit an authentication profile. | Yes | Yes | Yes |
| Authentication Sequence | Controls access to the **Authentication Sequence** node. If you disable this privilege, the administrator will not see the **Authentication Sequence** node or be able to create or edit an authentication sequence.<br>If you set this privilege to read only, the administrator can view the **Authentication Profile** information but cannot create or edit an authentication sequence. | Yes | Yes | Yes |
| Access Domain | Controls access to the **Access Domain** node. If you disable this privilege, the administrator will not see the **Access Domain** node or be able to create or edit an access domain.<br>If you set this privilege to read only, the administrator can view the **Access Domain** information but cannot create or edit an access domain. | Yes | Yes | Yes |
| Scheduled Log Export | Controls access to the **Scheduled Log Export** node. If you disable this privilege, the administrator will not see the **Scheduled Log Export** node or be able schedule exports of logs and save them to a File Transfer Protocol (FTP) server in CSV format or use Secure Copy (SCP) to securely transfer data between the device and a remote host.<br>If you set this privilege to read only, the administrator can view the **Scheduled Log Export Profile** information but cannot schedule the export of logs. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Software | Controls access to the **Software** node. If you disable this privilege, the administrator will not see the **Software** node or view the latest versions of the PAN-OS software available from Palo Alto Networks, read the release notes for each version, and select a release to download and install.<br>If you set this privilege to read only, the administrator can view the **Software** information but cannot download or install software. | Yes | Yes | Yes |
| GlobalProtect Client | Controls access to the **GlobalProtect Client** node. If you disable this privilege, the administrator will not see the **GlobalProtect Client** node or view available GlobalProtect releases, download the code or activate the GlobalProtect agent.<br>If you set this privilege to read only, the administrator can view the available **GlobalProtect Client** releases but cannot download or install the agent software. | Yes | Yes | Yes |
| Dynamic Updates | Controls access to the **Dynamic Updates** node. If you disable this privilege, the administrator will not see the **Dynamic Updates** node or be able to view the latest updates, read the release notes for each update, or select an update to upload and install.<br>If you set this privilege to read only, the administrator can view the available **Dynamic Updates** releases, read the release notes but cannot upload or install the software. | Yes | Yes | Yes |
| Licenses | Controls access to the **Licenses** node. If you disable this privilege, the administrator will not see the **Licenses** node or be able to view the licenses installed or activate licenses.<br>If you set this privilege to read only, the administrator can view the installed **Licenses**, but cannot perform license management functions. | Yes | Yes | Yes |
| Support | Controls access to the **Support** node. If you disable this privilege, the administrator will not see the **Support** node or be able to access product and security alerts from Palo Alto Networks or generate tech support or stats dump files.<br>If you set this privilege to read only, the administrator can view the **Support** node and access product and security alerts but cannot generate tech support or stats dump files. | Yes | Yes | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Master Key and Diagnostics | Controls access to the **Master Key and Diagnostics** node. If you disable this privilege, the administrator will not see the **Master Key and Diagnostics** node or be able to specify a master key to encrypt private keys on the firewall.<br>If you set this privilege to read only, the administrator can view the **Master Key and Diagnostics** node and view information about master keys that have been specified but cannot add or edit a new master key configuration. | Yes | Yes | Yes |

## Define User Privacy Settings in the Admin Role Profile

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Privacy | Sets the default state to enable or disable for all of the privacy settings described below. | Yes | N/A | Yes |
| Show Full IP addresses | When set to disable, full IP addresses obtained by traffic running through the Palo Alto firewall are not shown in logs or reports. In place of the IP addresses that are normally displayed, the relevant subnet is displayed.<br>Scheduled reports that are displayed in the interface through **Monitor > Reports** and reports that are sent via scheduled emails will still display full IP addresses. Because of this exception, we recommend that the following settings within the **Monitor** tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler. | Yes | N/A | Yes |
| Show User Names in Logs and Reports | When set to disable, user names obtained by traffic running through the Palo Alto Networks firewall are not shown in logs or reports. Columns where the user names would normally be displayed are empty.<br>Scheduled reports that are displayed in the interface through **Monitor > Reports** or reports that are sent via the email scheduler will still display user names. Because of this exception, we recommend that the following settings within the Monitor tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler. | Yes | N/A | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| View Pcap Files | When set to disable, packet capture files that are normally available within the Traffic, Threat and Data Filtering logs are not displayed. | Yes | N/A | Yes |

## Restrict Admin Access to Commit Functions

**Restrict User Access Using the Commit Setting**

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Commit | When set to disable, an admin cannot commit any changes to a configuration. | Yes | N/A | Yes |

## Provide Granular Access to Global Settings

**Restrict User Access Using the Global Settings**

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Global | Sets the default state to enable or disable for all of the global settings described below. In effect, this setting is only for System Alarms at this time. | Yes | N/A | Yes |
| System Alarms | When set to disable, an admin cannot view or acknowledge alarms that are generated. | Yes | N/A | Yes |

## Provide Granular Access to the Panorama Tab

The following table lists the **Panorama** tab access levels and the custom Panorama administrator roles for which they are available. Firewall administrators cannot access any of these privileges.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Setup | Specifies whether the administrator can view or edit Panorama setup information, such as **Management**, **Operations**, **Services**, **WildFire**, or **HSM**.<br>If you set the privilege to read only, the administrator can see the information but cannot edit it.<br>If you disable this privilege, the administrator cannot see or edit the information. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Templates | Specifies whether the administrator can view, edit, add, or delete templates.<br>If you set the privilege to read only, the administrator can see template configurations but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage template configurations. | Panorama: Yes<br>Device Group/Template: Yes | Yes<br>(No for Device Group and Template admins) | Yes | Yes |
| Config Audit | Specifies whether the administrator can run Panorama configuration audits. If you disable this privilege, the administrator cannot run Panorama configuration audits. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Managed Devices | Specifies whether the administrator can view, add, edit, tag, or delete firewalls as managed devices, and install software or content updates on them.<br>If you set this privilege to read only, the administrator can see managed firewalls but cannot add, delete, tag, or install updates on them.<br>If you disable this privilege, the administrator cannot view, add, edit, tag, delete, or install updates on managed firewalls.<br>This privilege applies only to the **Panorama > Managed Devices** page. An administrator with Device Deployment privileges can still use the **Panorama > Device Deployment** pages to install updates on managed firewalls. | Panorama: Yes<br>Device Group/Template: Yes | Yes<br>(No for Device Group and Template roles) | Yes | Yes |
| Device Groups | Specifies whether the administrator can view, edit, add, or delete device groups.<br>If you set this privilege to read only, the administrator can see device group configurations but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage device group configurations. | Panorama: Yes<br>Device Group/Template: Yes | Yes<br>(No for Device Group and Template roles) | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Managed Collectors | Specifies whether the administrator can view, edit, add, or delete Managed Collectors.<br>If you set this privilege to read only, the administrator can see Managed Collector configurations but cannot manage them.<br>If you disable this privilege, the administrator cannot view, edit, add, or delete Managed Collector configurations.<br>This privilege applies only to the **Panorama > Managed Collectors** page. An administrator with Device Deployment privileges can still use the **Panorama > Device Deployment** pages to install updates on Log Collectors. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Collector Groups | Specifies whether the administrator can view, edit, add, or delete Collector Groups.<br>If you set this privilege to read only, the administrator can see Collector Groups but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage Collector Groups. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Admin Roles | Specifies whether the administrator can view Panorama administrator roles.<br>You cannot enable full access to this function: just read only access. (Only Panorama administrators with a dynamic role can add, edit, or delete custom Panorama roles.) With read only access, the administrator can see Panorama administrator role configurations but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage Panorama administrator roles. | Panorama: Yes<br>Device Group/Template: No | No | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Administrators | Specifies whether the administrator can view Panorama administrator account details.<br>You cannot enable full access to this function: just read only access. (Only Panorama administrators with a dynamic role can add, edit, or delete Panorama administrators.) With read only access, the administrator can see information about his or her own account but no other Panorama administrator accounts.<br>If you disable this privilege, the administrator cannot see information about any Panorama administrator account, including his or her own. | Panorama: Yes<br>Device Group/Template: No | No | Yes | Yes |
| High Availability | Specifies whether the administrator can view and manage high availability (HA) settings for the Panorama management server.<br>If you set this privilege to read only, the administrator can view HA configuration information for the Panorama management server but cannot manage the configuration.<br>If you disable this privilege, the administrator cannot see or manage HA configuration settings for the Panorama management server. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| VMware Service Manager | Specifies whether the administrator can view and edit VMware Service Manager settings.<br>If you set this privilege to read only, the administrator can see the settings but cannot perform any related configuration or operational procedures.<br>If you disable this privilege, the administrator cannot see the settings or perform any related configuration or operational procedures. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Certificate Management | Sets the default state, enabled or disabled, for all of the Panorama certificate management privileges. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Certificates | Specifies whether the administrator can view, edit, generate, delete, revoke, renew, or export certificates. This privilege also specifies whether the administrator can import or export HA keys.<br>If you set this privilege to read only, the administrator can see Panorama certificates but cannot manage the certificates or HA keys.<br>If you disable this privilege, the administrator cannot see or manage Panorama certificates or HA keys. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Certificate Profile | Specifies whether the administrator can view, add, edit, delete or clone Panorama certificate profiles.<br>If you set this privilege to read only, the administrator can see Panorama certificate profiles but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage Panorama certificate profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Log Settings | Sets the default state, enabled or disabled, for all the log setting privileges. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| System | Specifies whether the administrator can see and configure the settings that control the forwarding of system logs to external services (Syslog, email, or SNMP trap servers).<br>If you set this privilege to read only, the administrator can see the system log forwarding settings but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the settings.<br>On a Panorama M-100 appliance, this privilege pertains only to system logs that Panorama generates for itself. On a Panorama virtual appliance, this privilege applies to system logs that Panorama generates and to system logs that Panorama collects from firewalls. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Config | Specifies whether the administrator can see and configure the settings that control the forwarding of configuration logs to external services (Syslog, email, or SNMP trap servers).<br>If you set this privilege to read only, the administrator can see the configuration log forwarding settings but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the settings.<br> On a Panorama M-100 appliance, this privilege pertains only to configuration logs that Panorama generates for itself. On a Panorama virtual appliance, this privilege applies to configuration logs that Panorama generates and to configuration logs that Panorama collects from firewalls. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| HIP Match | Specifies whether the administrator can see and configure the settings that control the forwarding of HIP match logs from a Panorama virtual appliance to external services (Syslog, email, or SNMP trap servers).<br> The **Panorama > Collector Groups** page controls the forwarding of HIP match logs from a Panorama M-100 appliance. The **Device > Log Settings > HIP Match** page controls the forwarding of HIP match logs directly from firewalls to external services (without aggregation on Panorama).<br>If you set this privilege to read only, the administrator can see the forwarding settings of HIP match logs but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the settings. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Traffic | Specifies whether the administrator can see and configure the settings that control the forwarding of traffic logs from a Panorama virtual appliance to external services (Syslog, email, or SNMP trap servers).<br> The **Panorama > Collector Groups** page controls the forwarding of traffic logs from a Panorama M-100 appliance. The **Objects > Log Forwarding** page controls the forwarding of traffic logs directly from firewalls to external services (without aggregation on Panorama).<br>If you set this privilege to read only, the administrator can see the forwarding settings of traffic logs but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the settings. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Threat | Specifies whether the administrator can see and configure the settings that control the forwarding of threat logs from a Panorama virtual appliance to external services (Syslog, email, or SNMP trap servers).<br> The **Panorama > Collector Groups** page controls the forwarding of threat logs from a Panorama M-100 appliance. The **Objects > Log Forwarding** page controls the forwarding of threat logs directly from firewalls to external services (without aggregation on Panorama).<br>If you set this privilege to read only, the administrator can see the forwarding settings of threat logs but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the settings. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Wildfire | Specifies whether the administrator can see and configure the settings that control the forwarding of WildFire logs from a Panorama virtual appliance to external services (Syslog, email, or SNMP trap servers).<br> The **Panorama > Collector Groups** page controls the forwarding of WildFire logs from a Panorama M-100 appliance. The **Objects > Log Forwarding** page controls the forwarding of WildFire logs directly from firewalls to external services (without aggregation on Panorama).<br>If you set this privilege to read only, the administrator can see the forwarding settings of WildFire logs but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the settings. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Server Profiles | Sets the default state, enabled or disabled, for all the server profile privileges.<br> These privileges pertain only to the server profiles that are used for forwarding logs that Panorama generates or collects from firewalls and the server profiles that are used for authenticating Panorama administrators. The **Device > Server Profiles** pages control the server profiles that are used for forwarding logs directly from firewalls to external services (without aggregation on Panorama) and for authenticating firewall administrators. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| SNMP Trap | Specifies whether the administrator can see and configure SNMP trap server profiles.<br>If you set this privilege to read only, the administrator can see SNMP trap server profiles but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage SNMP trap server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Syslog | Specifies whether the administrator can see and configure Syslog server profiles.<br>If you set this privilege to read only, the administrator can see Syslog server profiles but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage Syslog server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Email | Specifies whether the administrator can see and configure email server profiles.<br>If you set this privilege to read only, the administrator can see email server profiles but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage email server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| RADIUS | Specifies whether the administrator can see and configure the RADIUS server profiles that are used to authenticate Panorama administrators.<br>If you set this privilege to read only, the administrator can see the RADIUS server profiles but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the RADIUS server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| LDAP | Specifies whether the administrator can see and configure the LDAP server profiles that are used to authenticate Panorama administrators.<br>If you set this privilege to read only, the administrator can see the LDAP server profiles but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the LDAP server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Kerberos | Specifies whether the administrator can see and configure the Kerberos server profiles that are used to authenticate Panorama administrators.<br>If you set this privilege to read only, the administrator can see the Kerberos server profiles but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the Kerberos server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Authentication Profile | Specifies whether the administrator can view, add, edit, delete, or clone authentication profiles for Panorama administrators.<br>If you set this privilege to read only, the administrator can view Panorama authentication profiles but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage Panorama authentication profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Authentication Sequence | Specifies whether the administrator can view, add, edit, delete, or clone authentication sequences for Panorama administrators.<br>If you set this privilege to read only, the administrator can view Panorama authentication sequences but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage Panorama authentication sequences. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Access Domain | Specifies whether the administrator can view, add, edit, delete, or clone access domains for Panorama administrators.<br>If you set this privilege to read only, the administrator can view Panorama access domains but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage Panorama access domains. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Scheduled Config Export | Specifies whether the administrator can view, add, edit, delete, or clone scheduled Panorama configuration exports.<br>If you set this privilege to read only, the administrator can view the scheduled exports but cannot manage them.<br>If you disable this privilege, the administrator cannot see or manage the scheduled exports. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Software | Specifies whether the administrator can: view information about Panorama software updates; download, upload, or install the updates; and view the associated release notes.<br>If you set this privilege to read only, the administrator can view information about Panorama software updates and view the associated release notes but cannot perform any related operations.<br>If you disable this privilege, the administrator cannot see Panorama software updates, see the associated release notes, or perform any related operations.<br>This privilege pertains only to software installed on a Panorama management server. The **Panorama > Device Deployment > Software** page controls access to PAN-OS software deployed on firewalls and Panorama software deployed on dedicated Log Collectors. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Dynamic Updates | Specifies whether the administrator can: view information about Panorama content updates (for example, WildFire updates); download, upload, install, or revert the updates; and view the associated release notes.<br>If you set this privilege to read only, the administrator can view information about Panorama content updates and view the associated release notes but cannot perform any related operations.<br>If you disable this privilege, the administrator cannot see Panorama content updates, see the associated release notes, or perform any related operations.<br> This privilege pertains only to content updates installed on a Panorama management server. The **Panorama > Device Deployment > Dynamic Updates** page controls access to content updates deployed on firewalls and dedicated Log Collectors. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Support | Specifies whether the administrator can: view Panorama support license information, product alerts, and security alerts; activate a support license, generate Tech Support files, and manage cases<br>If you set this privilege to read only, the administrator can view Panorama support information, product alerts, and security alerts, but cannot activate a support license, generate Tech Support files, or manage cases.<br>If you disable this privilege, the administrator cannot: see Panorama support information, product alerts, or security alerts; activate a support license, generate Tech Support files, or manage cases. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Device Deployment | Sets the default state, enabled or disabled, for all the device deployment privileges.<br> These privilege pertain only to software and content updates that Panorama administrators deploy on firewalls and dedicated Log Collectors. The **Panorama > Software** and **Panorama > Dynamic Updates** pages control the software and content updates installed on a Panorama management server. | Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Software | Specifies whether the administrator can: view information about the software updates installed on firewalls and Log Collectors; download, upload, or install the updates; and view the associated release notes.<br>If you set this privilege to read only, the administrator can see information about the software updates and view the associated release notes but cannot deploy the updates to firewalls or dedicated Log Collectors.<br>If you disable this privilege, the administrator cannot see information about the software updates, see the associated release notes, or deploy the updates to firewalls or dedicated Log Collectors. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |
| SSL VPN Client | Specifies whether the administrator can: view information about SSL VPN client software updates on firewalls; download, upload, or activate the updates; and view the associated release notes.<br>If you set this privilege to read only, the administrator can see information about SSL VPN client software updates and view the associated release notes but cannot activate the updates on firewalls.<br>If you disable this privilege, the administrator cannot see information about SSL VPN client software updates, see the associated release notes, or activate the updates on firewalls. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| GlobalProtect Client | Specifies whether the administrator can: view information about GlobalProtect client software updates on firewalls; download, upload, or activate the updates; and view the associated release notes.<br>If you set this privilege to read only, the administrator can see information about GlobalProtect client software updates and view the associated release notes but cannot activate the updates on firewalls.<br>If you disable this privilege, the administrator cannot see information about GlobalProtect client software updates, see the associated release notes, or activate the updates on firewalls. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |
| Dynamic Updates | Specifies whether the administrator can: view information about the content updates (for example, Applications updates) installed on firewalls and dedicated Log Collectors; download, upload, or install the updates; and view the associated release notes.<br>If you set this privilege to read only, the administrator can see information about the content updates and view the associated release notes but cannot deploy the updates to firewalls or dedicated Log Collectors.<br>If you disable this privilege, the administrator cannot see information about the content updates, see the associated release notes, or deploy the updates to firewalls or dedicated Log Collectors. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |
| Licenses | Specifies whether the administrator can view, refresh, and activate firewall licenses.<br>If you set this privilege to read only, the administrator can view firewall licenses but cannot refresh or activate those licenses.<br>If you disable this privilege, the administrator cannot view, refresh, or activate firewall licenses. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Master Key and Diagnostics | Specifies whether the administrator can view and configure a master key by which to encrypt private keys on Panorama.<br>If you set this privilege to read only, the administrator can view the Panorama master key configuration but cannot change it.<br>If you disable this privilege, the administrator cannot see or edit the Panorama master key configuration. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

## Panorama Web Interface Access

The custom Panorama administrator roles allow you to define access to the options on Panorama and the ability to only allow access to Device Groups and Templates (**Policies**, **Objects**, **Network**, **Device** tabs).

The admin roles you can create are **Panorama** and **Device Group and Template**. You cannot assign CLI access privileges to a **Device Group and Template** admin role. If you assign superuser privileges for the CLI to a **Panorama** admin role, administrators with that role can access all features regardless of the web interface privileges you assign.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Dashboard | Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets. | Yes | No | Yes |
| ACC | Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Yes | No | Yes |
| Monitor | Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the admin can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab. | Yes | No | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Policies | Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the admin can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab. | Yes | No | Yes |
| Objects | Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab. | Yes | No | Yes |
| Network | Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the admin can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab. | Yes | No | Yes |
| Device | Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any device-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the admin can see, leave the **Device** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab.<br> You cannot enable access to the **Admin Roles** or **Administrators** nodes for a role-based administrator even if you enable full access to the **Device** tab. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Panorama | Controls access to the **Panorama** tab. If you disable this privilege, the administrator will not see the **Panorama** tab and will not have access to any Panorama-wide configuration information, such as Managed Devices, Managed Collectors, or Collector Groups.<br>For more granular control over what objects the admin can see, leave the **Panorama** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Panorama Tab. | Yes | No | Yes |

# Reference: Port Numbers Used by Palo Alto Networks Devices

The following tables list the ports that Palo Alto Networks devices use to communicate with each other, or with other services on the network.


- Ports Used for Management Functions
- Ports Used for HA
- Ports Used for Panorama
- Ports Used for User-ID


## Ports Used for Management Functions

| Destination Port | Protocol | Description |
|---|---|---|
| 22 | TCP | Used for communication from a client system to the firewall CLI interface. |
| 80 | TCP | The port the firewall listens on for Online Certificate Status Protocol (OCSP) updates when acting as an OCSP responder. |
| 123 | UDP | Port the firewall uses for NTP updates. |
| 443 | TCP | Used for communication from a client system to the firewall web interface. This is also the port the firewall and/or the User-ID agent listens on for VM Information source updates. For monitoring an AWS environment, this is the only port that is used. For monitoring a VMware vCenter/ESXi environment, the listening port defaults to 443, but it is configurable. |
| 162 | UDP | Port the firewall, Panorama, or the log collector use to send SNMP traps to a SNMP trap receiver.<br> This port does not need to be open on the Palo Alto Networks device. The SNMP trap receiver must be configured to listen on this port. |
| 161 | UDP | Port the firewall listens on for SNMP polling requests from the NMS. |
| 514<br>514<br>6514 | TCP<br>UDP<br>SSL | Port the firewall, Panorama, or the log collector use to send logs to a Syslog server and the ports the User-ID agent (on the firewall or on a Windows server) listens on for authentication Syslog messages for use with User-ID. |
| 2055 | UDP | Default port the firewall uses to send NetFlow records, but this is configurable. |
| 5008 | TCP | Port the GlobalProtect Mobile Security Manager listens on for HIP requests from the GlobalProtect gateways.<br>If you are using a third-party MDM system, you can configure the gateway to use a different port as required by the MDM vendor. |

 © Palo Alto Networks, Inc.

| Destination Port | Protocol | Description |
|---|---|---|
| 6081<br>6082 | TCP<br>TCP | Ports used for Captive Portal. |

## Ports Used for HA

Firewalls configured as High Availability (HA) peers must be able to communicate with each other to maintain state information (HA1 control link) and synchronize data (HA2 data link). In Active/Active HA deployments the peer firewalls must also forward packets to the HA peer that owns the session. The HA3 link is a Layer 2 (MAC-in-MAC) link and it does not support Layer 3 addressing or encryption.

| Destination Port | Protocol | Description |
|---|---|---|
| 28769<br>28260 | TCP<br>TCP | Used for the HA1 control link for clear text communication between the HA peer firewalls. The HA1 link is a Layer 3 link and requires an IP address. |
| 28 | TCP | Used for the HA1 control link for encrypted communication (SSH over TCP) between the HA peer firewalls. |
| 28770 | TCP | Listening port for HA1 backup links. |
| 28771 | TCP | Used for heartbeat backups. Palo Alto Networks recommends enabling heartbeat backup on the MGT interface if you use an in-band port for the HA1 or the HA1 backup links. |
| 99<br>29281 | IP<br>UDP | Used for the HA2 link to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between firewalls in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active device (Active/Passive) or active-primary (Active/Active) to the passive device (Active/Passive) or active-secondary (Active/Active). The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.<br>The HA data link can also be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets. |

## Ports Used for Panorama

| Destination Port | Protocol | Description |
|---|---|---|
| 22 | TCP | Used for communication from a client system to the Panorama CLI interface. |
| 443 | TCP | Used for communication from a client system to the Panorama web interface. |
| 3978 | TCP | Used for communication between Panorama and managed devices and/or managed log collectors as well as for communication between log collectors in a collector group as follows:<br>• For communication between Panorama and managed devices, this is a bi-directional connection on which the managed firewalls forward logs to Panorama and for Panorama to push configuration changes to the managed devices. Context switching commands are sent over the same connection.<br>• Log collectors use this destination port to forward logs to Panorama.<br>• For communication between the default log collector on a Panorama in Panorama mode, and for communicating with log collectors in a DLC architecture mode. |
| 28769 (5.1 and later)<br>28260 (5.0 and later)<br>49160 (5.0 and earlier) | TCP<br>TCP<br>TCP | Used for the HA connectivity and synchronization between Panorama HA peers using clear text communication. Communication can be initiated by either peer. |
| 28 | TCP | Used for the HA connectivity and synchronization between Panorama HA peers using encrypted communication (SSH over TCP). Communication can be initiated by either peer. |
| 28270 (6.0 and later)<br>49190 (5.1 and earlier) | TCP | Used for communication between log collectors in a collector group for log distribution. |
| 2049 | TCP | Used by the Panorama virtual appliance to write logs to the NFS datastore. |

 © Palo Alto Networks, Inc.

## Ports Used for User-ID

User-ID is a feature that enables mapping of user IP addresses to usernames and group memberships, enabling user- or group-based policy and visibility into user activity on your network (for example, to be able to quickly track down a user who may be the victim of a threat). To perform this mapping, the firewall, the User-ID agent (either installed on a Windows-based system or the PAN-OS integrated agent running on the firewall), and/or the Terminal Services agent must be able to connect to directory services on your network to perform Group Mapping and User Mapping. Additionally, if the agents are running on systems external to the firewall, they must be able to connect to the firewall to communicate the IP address to username mappings to the firewall. The following table lists the communication requirements for User-ID along with the port numbers required to establish connections.

| Destination Port | Protocol | Description |
|---|---|---|
| 389 | TCP | Port the firewall uses to connect to the LDAP server (plaintext or StartTLS) in order to Map Users to Groups. |
| 636 | TCP | Port the firewall uses to connect to the LDAP server (LDAP over SSL) in order to Map Users to Groups. Used for LDAP over SSL connections. |
| 514<br>514<br>6514 | TCP<br>UDP<br>SSL | Ports the User-ID agent (on the firewall or on a Windows server) listens on for authentication Syslog messages for use with User-ID. |
| 5007 | TCP | Port the firewall listens on for user mapping information from the User-ID or Terminal Server agent. The agent sends the IP address and username mapping along with a timestamp whenever it learns of a new or updated mapping. In addition, it connects to the firewall at regular intervals to refresh known mappings. |
| 5006 | TCP | Port the User-ID agent listens on for User-ID XML API requests. The source for this communication is typically the system running a script that invokes the API. |
| 88 | UDP | Port the User-ID agent uses to authenticate to a Kerberos server. |
| 1812 | UDP | Port the User-ID agent uses to authenticate to a RADIUS server. |
| 135 | TCP | Port the User-ID agent uses to establish TCP-based WMI connections with the Microsoft Remote Procedure Call (RPC) Endpoint Mapper. The Endpoint Mapper then assigns the agent a randomly assigned port in the 49152-65535 port range. The agent uses this connection to make RPC queries for Exchange Server or AD server security logs, session tables. This is also the port used to access Terminal Services.<br>The User-ID agent also uses this port to connect to client systems to perform WMI probes. |
| 139 | TCP | Port the User-ID agent uses to establish TCP-based NetBIOS connections to the AD server so that it can send RPC queries for security logs and session information.<br>The User-ID agent also uses this port to connect to client systems for NetBIOS probing (supported on the Windows-based User-ID agent only). |
| 445 | TCP | Port the User-ID agent uses to connect to the Active Directory (AD) using TCP-based SMB connections to the AD server for access to user logon information (print spooler and Net Logon). |

# Reset the Firewall to Factory Default Settings

Resetting the firewall to factory defaults will result in the loss of all configuration settings and logs.

| Reset the Firewall to Factory Default Settings | |
|---|---|
| Step 1 Set up a console connection to the firewall. | 1. Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1).<br>If your computer does not have a 9-pin serial port, use a USB-to-serial port connector.<br>2. Enter your login credentials.<br>3. Enter the following CLI command:<br>`debug system maintenance-mode`<br>The firewall will reboot in the maintenance mode. |
| Step 2 Reset the system to factory default settings. | 1. When the device reboots, press `Enter` to continue to the maintenance mode menu.<br> <br>2. Select `Factory Reset` and press `Enter`.<br>3. Select `Factory Reset` and press `Enter` again.<br>The firewall will reboot without any configuration settings. The default username and password to log in to the firewall is admin/admin.<br>To perform initial configuration on the firewall and to set up network connectivity, see Integrate the Firewall into Your Management Network. |

 © Palo Alto Networks, Inc.

# Certificate Management

The following topics describe the different keys and certificates that Palo Alto Networks devices use, and how to obtain and manage them:

- Keys and Certificates
- Certificate Revocation
- Certificate Deployment
- Set Up Verification for Certificate Revocation Status
- Configure the Master Key
- Obtain Certificates
- Configure a Certificate Profile
- Configure the Key Size for SSL Forward Proxy Server Certificates
- Revoke and Renew Certificates
- Secure Keys with a Hardware Security Module

# Keys and Certificates

To ensure trust between parties in a secure communication session, Palo Alto Networks devices use digital certificates. Each certificate contains a cryptographic key to encrypt plaintext or decrypt cyphertext. Each certificate also includes a digital signature to authenticate the identity of the issuer. The issuer must be in the list of trusted certificate authorities (CAs) of the authenticating party. Optionally, the authenticating party verifies the issuer did not revoke the certificate (see Certificate Revocation).

Palo Alto Networks devices use certificates in the following applications:

- User authentication for Captive Portal, GlobalProtect, Mobile Security Manager, and firewall/Panorama web interface access.
- Device authentication for GlobalProtect VPN (remote user-to-site or large scale).
- Device authentication for IPSec site-to-site VPN with Internet Key Exchange (IKE).
- Decrypting inbound and outbound SSL traffic. A firewall decrypts the traffic to apply security policies and rules, then re-encrypts it before forwarding the traffic to the final destination. For outbound traffic, the firewall acts as a forward proxy server, establishing an SSL/TLS connection to the destination server. To secure a connection between itself and the client, the firewall uses a *signing certificate* to automatically generate a copy of the destination server certificate.

The following table describes the keys and certificates that Palo Alto Networks devices use. As a best practice, use different keys/certificates for each usage.

**Table: Palo Alto Networks Device Keys/Certificates**

| Key/Certificate Usage | Description |
|---|---|
| Administrative Access | Secure access to device administration interfaces (HTTPS access to the web interface) requires a server certificate for the MGT interface (or a designated interface on the dataplane if the device does not use MGT) and, optionally, a certificate to authenticate the administrator. |
| Captive Portal | In deployments where Captive Portal identifies users who access HTTPS resources, designate a server certificate for the Captive Portal interface. If you configure Captive Portal to use certificates (instead of, or in addition to, username/password credentials) for user identification, designate a user certificate also. For more information on Captive Portal, see Map IP Addresses to User Names Using Captive Portal. |
| Forward Trust | For outbound SSL/TLS traffic, if a firewall acting as a forward proxy trusts the CA that signed the certificate of the destination server, the firewall uses the forward trust CA certificate to generate a copy of the destination server certificate to present to the client. To set the key size, see Configure the Key Size for SSL Forward Proxy Server Certificates. For added security, store the key on a hardware security module (for details, see Secure Keys with a Hardware Security Module). |
| Forward Untrust | For outbound SSL/TLS traffic, if a firewall acting as a forward proxy does not trust the CA that signed the certificate of the destination server, the firewall uses the forward untrust CA certificate to generate a copy of the destination server certificate to present to the client. |

 © Palo Alto Networks, Inc.

| Key/Certificate Usage | Description |
|---|---|
| SSL Inbound Inspection | The keys that decrypt inbound SSL/TLS traffic for inspection and policy enforcement. For this application, import onto the firewall a private key for each server that is subject to SSL/TLS inbound inspection. See Configure SSL Inbound Inspection. |
| SSL Exclude Certificate | Certificates for servers to exclude from SSL/TLS decryption. For example, if you enable SSL decryption but your network includes servers for which the firewall should not decrypt traffic (for example, web services for your HR systems), import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Configure Decryption Exceptions. |
| GlobalProtect | All interaction among GlobalProtect components occurs over SSL/TLS connections. Therefore, as part of the GlobalProtect deployment, deploy server certificates for all GlobalProtect portals, gateways, and Mobile Security Managers. Optionally, deploy certificates for authenticating users also.<br>Note that the GlobalProtect Large Scale VPN (LSVPN) feature requires a CA signing certificate. |
| Site-to-Site VPNs (IKE) | In a site-to-site IPSec VPN deployment, peer devices use Internet Key Exchange (IKE) gateways to establish a secure channel. IKE gateways use certificates or preshared keys to authenticate the peers to each other. You configure and assign the certificates or keys when defining an IKE gateway on a firewall. See Site-to-Site VPN Overview. |
| Master Key | The firewall uses a master key to encrypt all private keys and passwords. If your network requires a secure location for storing private keys, you can use an encryption (wrapping) key stored on a hardware security module (HSM) to encrypt the master key. For details, see Encrypt a Master Key Using an HSM. |
| Secure Syslog | The certificate to enable secure connections between the firewall and a syslog server. See Configure the Firewall to Authenticate to the Syslog Server. |
| Trusted Root CA | The designation for a root certificate issued by a CA that the firewall trusts. The firewall can use a self-signed root CA certificate to automatically issue certificates for other applications (for example, SSL Forward Proxy).<br>Also, if a firewall must establish secure connections with other firewalls, the root CA that issues their certificates must be in the list of trusted root CAs on the firewall. |

© Palo Alto Networks, Inc.

# Certificate Revocation

Palo Alto Networks devices use digital certificates to ensure trust between parties in a secure communication session. Configuring a device to check the revocation status of certificates provides additional security. A party that presents a revoked certificate is not trustworthy. When a certificate is part of a chain, the device checks the status of every certificate in the chain except the root CA certificate, for which the device cannot verify revocation status.

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority that issued the certificate must revoke it.

Palo Alto Networks devices support the following methods for verifying certificate revocation status. If you configure both, the devices first try the OCSP method; if the OCSP server is unavailable, the devices use the CRL method.

- Certificate Revocation List (CRL)
- Online Certificate Status Protocol (OCSP)

In PAN-OS, certificate revocation status verification is an optional feature. It is a best practice to enable it for certificate profiles, which define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPsec VPN, and firewall/Panorama web interface access.

## Certificate Revocation List (CRL)

Each certificate authority (CA) periodically issues a certificate revocation list (CRL) to a public repository. The CRL identifies revoked certificates by serial number. After the CA revokes a certificate, the next CRL update will include the serial number of that certificate.

The Palo Alto Networks firewall downloads and caches the last-issued CRL for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the CRL for the issuing CA. Also, the cache only stores a CRL until it expires.

The firewall supports CRLs only in Distinguished Encoding Rules (DER) format. If the firewall downloads a CRL in any other format—for example, Privacy Enhanced Mail (PEM) format—any revocation verification process that uses that CRL will fail when a user performs an activity that triggers the process (for example, sending outbound SSL data). The firewall will generate a system log for the verification failure. If the verification was for an SSL certificate, the firewall will also display the SSL Certificate Errors Notify response page to the user.

To use CRLs for verifying the revocation status of certificates used for the decryption of inbound and outbound SSL/TLS traffic, see Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

To use CRLs for verifying the revocation status of certificates that authenticate users and devices, configure a certificate profile and assign it to the interfaces that are specific to the application: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, or firewall/Panorama web interface access. For details, see Configure Revocation Status Verification of Certificates Used for User/Device Authentication.

 © Palo Alto Networks, Inc.

## Online Certificate Status Protocol (OCSP)

When establishing an SSL/TLS session, clients can use Online Certificate Status Protocol (OCSP) to check the revocation status of the authentication certificate. The authenticating client sends a request containing the serial number of the certificate to the OCSP responder (server). The responder searches the database of the certificate authority (CA) that issued the certificate and returns a response containing the status (*good*, *revoked* or *unknown*) to the client. The advantage of the OCSP method is that it can verify status in real-time, instead of depending on the issue frequency (hourly, daily, or weekly) of CRLs.

The Palo Alto Networks firewall downloads and caches OCSP status information for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the OCSP information for the issuing CA. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall as an OCSP responder (see Configure an OCSP Responder).

To use OCSP for verifying the revocation status of certificates when the firewall functions as an SSL forward proxy, perform the steps under Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

The following applications use certificates to authenticate users and/or devices: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, and firewall/Panorama web interface access. To use OCSP for verifying the revocation status of the certificates:

- Configure an OCSP responder.
- Enable the HTTP OCSP service on the firewall.
- Create or obtain a certificate for each application.
- Configure a certificate profile for each application.
- Assign the certificate profile to the relevant application.

To cover situations where the OCSP responder is unavailable, configure CRL as a fall-back method. For details, see Configure Revocation Status Verification of Certificates Used for User/Device Authentication.

# Certificate Deployment

The basic approaches to deploy certificates for Palo Alto Networks devices are:

- **Obtain certificates from a trusted third-party CA**—The benefit of obtaining a certificate from a trusted third-party certificate authority (CA) such as VeriSign or GoDaddy is that end clients will already trust the certificate because common browsers include root CA certificates from well-known CAs in their trusted root certificate stores. Therefore, for applications that require end clients to establish secure connections with a Palo Alto Network device, purchase a certificate from a CA that the end clients trust to avoid having to pre-deploy root CA certificates to the end clients. (Some such applications are a GlobalProtect portal or GlobalProtect Mobile Security Manager.) However, note that most third-party CAs cannot issue signing certificates. Therefore, this type of certificate is not appropriate for applications (for example, SSL/TLS decryption and large-scale VPN) that require the firewall to issue certificates. See Obtain a Certificate from an External CA.

- **Obtain certificates from an enterprise CA**—Enterprises that have their own internal CA can use it to issue certificates for firewall applications and import them onto the firewall. The benefit is that end clients probably already trust the enterprise CA. You can either generate the needed certificates and import them onto the firewall, or generate a certificate signing request (CSR) on the firewall and send it to the enterprise CA for signing. The benefit of this method is that the private key does not leave the firewall. An enterprise CA can also issue a signing certificate, which the firewall uses to automatically generate certificates (for example, for GlobalProtect large-scale VPN or sites requiring SSL/TLS decryption). See Import a Certificate and Private Key.

- **Generate self-signed certificates**—You can Create a Self-Signed Root CA Certificate on the firewall and use it to automatically issue certificates for other firewall applications. Note that if you use this method to generate certificates for an application that requires an end client to trust the certificate, end users will see a certificate error because the root CA certificate is not in their trusted root certificate store. To prevent this, deploy the self-signed root CA certificate to all end user systems. You can deploy the certificates manually or use a centralized deployment method such as an Active Directory Group Policy Object (GPO).

 © Palo Alto Networks, Inc.

# Set Up Verification for Certificate Revocation Status

To verify the revocation status of certificates, the firewall uses Online Certificate Status Protocol (OCSP) and/or certificate revocation lists (CRLs). For details on these methods, see Certificate Revocation If you configure both methods, the firewall first tries OCSP and only falls back to the CRL method if the OCSP responder is unavailable. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall to function as the OCSP responder.

The following topics describe how to configure the firewall to verify certificate revocation status:

- Configure an OCSP Responder
- Configure Revocation Status Verification of Certificates Used for User/Device Authentication
- Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

## Configure an OCSP Responder

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of certificates, you must configure the firewall to access an OCSP responder (server). The entity that manages the OCSP responder can be a third-party certificate authority (CA) or, if your enterprise has its own public key infrastructure (PKI), the firewall itself. For details on OCSP, see Certificate Revocation

| **Configure an OCSP Responder** | |
|---|---|
| Step 1 Define an OCSP responder. | 1. In a firewall, select **Device > Certificate Management > OCSP Responder** and click **Add**.<br>In Panorama, select **Device > Certificate Management > OCSP Responder**, select a **Template** and click **Add**.<br>2. Enter a **Name** to identify the responder (up to 31 characters). The name is case-sensitive. It must be unique and use only letters, numbers, spaces, hyphens, and underscores.<br>3. If the firewall supports multiple virtual systems, the dialog displays a **Location** drop-down. Select the virtual system where the responder will be available or select **Shared** to enable availability on all the virtual systems.<br>4. In the **Host Name** field, enter the host name (recommended) or IP address of the OCSP responder. From this value, PAN-OS automatically derives a URL and adds it to the certificate being verified.<br>If you configure the firewall itself as an OCSP responder, the host name must resolve to an IP address in the interface that the firewall uses for OCSP services (specified in Step 3).<br>5. Click **OK**. |

| Configure an OCSP Responder | | |
|---|---|---|
| Step 2 | Enable OCSP communication on the firewall. | 1. In a firewall, select **Device > Setup > Management**.<br>In Panorama, select **Device > Setup > Management** and select a **Template**.<br>2. In the Management Interface Settings section, edit to select the **HTTP OCSP** check box, then click **OK**. |
| Step 3 | Optionally, to configure the firewall itself as an OCSP responder, add an Interface Management Profile to the interface used for OCSP services. | 1. Select **Network > Network Profiles > Interface Mgmt**.<br>2. Click **Add** to create a new profile or click the name of an existing profile.<br>3. Select the **HTTP OCSP** check box and click **OK**.<br>4. Select **Network > Interfaces** and click the name of the interface that the firewall will use for OCSP services. The OCSP **Host Name** specified in Step 1 must resolve to an IP address in this interface.<br>5. Select **Advanced > Other info** and select the Interface Management Profile you configured.<br>6. Click **OK** and **Commit**. |

## Configure Revocation Status Verification of Certificates Used for User/Device Authentication

The firewall uses certificates to authenticate users and devices for such applications as Captive Portal, GlobalProtect, site-to-site IPSec VPN, and firewall/Panorama web interface access. To improve security, it is a best practice to configure the firewall to verify the revocation status of certificates that it uses for device/user authentication.

| Configure Revocation Status Verification of Certificates Used for User/Device Authentication | | |
|---|---|---|
| Step 1 | Configure a Certificate Profile for each application. | Assign one or more root CA certificates to the profile and select how the firewall verifies certificate revocation status. The common name (FQDN or IP address) of a certificate must match an interface to which you apply the profile in Step 2.<br>For details on the certificates that various applications use, see Keys and Certificates |
| Step 2 | Assign the certificate profiles to the relevant applications. | The steps to assign a certificate profile depend on the application that requires it. |

 © Palo Alto Networks, Inc.

# Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

The firewall decrypts inbound and outbound SSL/TLS traffic to apply security policies and rules, then re-encrypts the traffic before forwarding it. (For details, see SSL Inbound Inspection and SSL Forward Proxy.) You can configure the firewall to verify the revocation status of certificates used for decryption as follows.

Enabling revocation status verification for SSL/TLS decryption certificates will add time to the process of establishing the session. The first attempt to access a site might fail if the verification does not finish before the session times out. For these reasons, verification is disabled by default.

| **Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption** | | |
|---|---|---|
| Step 1 | Access the Decryption Certificate Revocation Settings page. | In a firewall, select **Device > Setup > Session** and, in the Session Features section, select **Decryption Certificate Revocation Settings**.<br>In Panorama, select **Device > Setup > Session**, select a **Template** and, in the Session Features section, select **Decryption Certificate Revocation Settings**. |
| Step 2 | Define the service-specific timeout intervals for revocation status requests. | Perform one or both of the following steps, depending on whether the firewall will use Online Certificate Status Protocol (OCSP) or the Certificate Revocation List (CRL) method to verify the revocation status of certificates. If the firewall will use both, it first tries OCSP; if the OCSP responder is unavailable, the firewall then tries the CRL method.<br>1. In the CRL section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the CRL service.<br>2. In the OCSP section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the OCSP responder.<br>Depending on the **Certificate Status Timeout** value you specify in Step 3, the firewall might register a timeout before either or both of the **Receive Timeout** intervals pass. |

| Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption (Continued) | | |
|---|---|---|
| Step 3 | Define the total timeout interval for revocation status requests. | Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies the session-blocking logic you optionally define in Step 4. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:<br>• If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.<br>• If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.<br>• If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value. |
| Step 4 | Define the blocking behavior for *unknown* certificate status or a revocation status request timeout. | If you want the firewall to block SSL/TLS sessions when the OCSP or CRL service returns a certificate revocation status of *unknown*, select the **Block Session With Unknown Certificate Status** check box. Otherwise, the firewall proceeds with the session.<br>If you want the firewall to block SSL/TLS sessions after it registers a request timeout, select the **Block Session On Certificate Status Check Timeout** check box. Otherwise, the firewall proceeds with the session. |
| Step 5 | Save and apply your entries. | Click **OK** and **Commit**. |


© Palo Alto Networks, Inc.

# Configure the Master Key

Every firewall has a default master key that encrypts private keys and other secrets (such as passwords and shared keys). The private keys authenticate users when they access administrative interfaces on the firewall. As a best practice to safeguard the keys, configure the master key on each firewall to be unique and periodically change it. For added security, use a *wrapping key* stored on a hardware security module (HSM) to encrypt the master key. For details, see Encrypt a Master Key Using an HSM.

In a high availability (HA) configuration, ensure both devices in the pair use the same master key to encrypt private keys and certificates. If the master keys differ, HA configuration synchronization will not work properly.

When you export a firewall configuration, the master key encrypts the passwords of users managed on external servers. For locally managed users, the firewall hashes the passwords but the master key does not encrypt them.

**Configure a Master Key**

1. In a firewall, select **Device > Master Key and Diagnostics** and, in the Master Key section, click the Edit icon.
   In Panorama, select **Panorama > Master Key and Diagnostics** and, in the Master Key section, click the Edit icon.
2. Enter the **Current Master Key** if one exists.
3. Define a new **New Master Key** and then **Confirm New Master Key**. The key must contain exactly 16 characters.
4. (Optional) To specify the master key **Life Time**, enter the number of **Days** and/or **Hours** after which the key will expire. If you set a life time, create a new master key before the old key expires.
5. (Optional) If you set a key life time, enter a **Time for Reminder** that specifies the number of **Days** and **Hours** preceding master key expiration when the firewall emails you a reminder.
6. (Optional) Select whether to use an **HSM** to encrypt the master key. For details, see Encrypt a Master Key Using an HSM.
7. Click **OK** and **Commit**.

# Obtain Certificates


- Create a Self-Signed Root CA Certificate
- Generate a Certificate on the Firewall
- Import a Certificate and Private Key
- Obtain a Certificate from an External CA



© Palo Alto Networks, Inc.

## Create a Self-Signed Root CA Certificate

A self-signed root certificate authority (CA) certificate is the top-most certificate in a certificate chain. A firewall can use this certificate to automatically issue certificates for other uses. For example, the firewall issues certificates for SSL/TLS decryption and for satellite devices in a GlobalProtect large-scale VPN.

When establishing a secure connection with the firewall, the remote client must trust the root CA that issued the certificate. Otherwise, the client browser will display a warning that the certificate is invalid and might (depending on security settings) block the connection. To prevent this, after generating the self-signed root CA certificate, import it into the client systems.

**Generate a Self-signed Root CA Certificate**

1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.
   In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.
2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in Step 6.
3. Click **Generate**.
4. Enter a **Certificate Name**, such as *GlobalProtect_CA*. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.
5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.
6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the device supports multiple virtual systems.
7. Leave the **Signed By** field blank to designate the certificate as self-signed.
8. Select the **Certificate Authority** check box.
9. Do *not* select an **OCSP Responder**. Certificate revocation status verification does not apply to root CA certificates.
10. Click **Generate** and **Commit**.

## Generate a Certificate on the Firewall

The firewall uses certificates to authenticate clients, servers, users and devices in several applications, including SSL/TLS decryption, Captive Portal, GlobalProtect, site-to-site IPSec VPN, and firewall/Panorama web interface access. Generate certificates for each usage. For details on application-specific certificates, see Keys and Certificates

To generate a certificate, you must first create or import a root CA certificate to sign it. For details, see Create a Self-Signed Root CA Certificate or Import a Certificate and Private Key.

To use Online Certificate Status Protocol (OCSP) for verifying certificate revocation status, Configure an OCSP Responder *before* generating the certificate. For details on status verification, see Certificate Revocation

**Generate a Certificate on the Firewall**

1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.
   In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.
2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in Step 6.
3. Click **Generate**.
4. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.
5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.
6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the device supports multiple virtual systems.
7. In the **Signed By** field, select the root CA certificate that will issue the certificate.
8. If applicable, select an **OCSP Responder**.
9. (Optional) Define the **Cryptographic Settings** as necessary to create a certificate that will work with the devices that must authenticate to it. The default and recommended key size (**Number of Bits**) is 2048 bits. The default and recommended encryption algorithm (**Digest**) is SHA256.
10. (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.
    If you add a **Host Name** (DNS name) attribute, it is a best practice for it to match the **Common Name**. The host name populates the Subject Alternative Name field of the certificate.
11. Click **Generate** and, in the Device Certificates tab, click the certificate **Name**.
    Regardless of the time zone on the firewall, it always displays the corresponding Greenwich Mean Time (GMT) for certificate validity and expiration dates/times.
12. Select the check boxes that correspond to the intended use of the certificate on the firewall. For example, if the firewall will use this certificate to authenticate user access to its web interface, select the **Certificate for Secure Web GUI** check box.
13. Click **OK** and **Commit**.

 © Palo Alto Networks, Inc.

## Import a Certificate and Private Key

If your enterprise has its own public key infrastructure (PKI), you can import a certificate and private key into the firewall from your enterprise certificate authority (CA). Enterprise CA certificates (unlike most certificates purchased from a trusted, third-party CA) can automatically issue CA certificates for applications such as SSL/TLS decryption or large-scale VPN.

Instead of importing a self-signed root CA certificate into all the client systems, it is a best practice to import a certificate from the enterprise CA because the clients will already have a trust relationship with the enterprise CA, which simplifies the deployment.

If the certificate you will import is part of a certificate chain, it is a best practice to import the entire chain.

**Import a Certificate and Private Key**

1. From the enterprise CA, export the certificate and private key that the firewall will use for authentication.

   When exporting a private key, you must enter a passphrase to encrypt the key for transport. Ensure the management system can access the certificate and key files. When importing the key onto the firewall, you must enter the same passphrase to decrypt it.

2. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.

   In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.

3. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in Step 6.

4. Click **Import**.

5. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.

6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the device supports multiple virtual systems.

7. Enter the path and name of the **Certificate File** received from the CA, or **Browse** to find the file.

8. Select a **File Format**:

   - **Encrypted Private Key and Certificate (PKCS12)**—This is the default and most common format, in which the key and certificate are in a single container (**Certificate File**). If a hardware security module (HSM) will store the private key for this certificate, select the **Private key resides on Hardware Security Module** check box.
   - **Base64 Encoded Certificate (PEM)**—You must import the key separately from the certificate. If a hardware security module (HSM) stores the private key for this certificate, select the **Private key resides on Hardware Security Module** check box and skip Step 9. Otherwise, select the **Import Private Key** check box, enter the **Key File** or **Browse** to it, then perform Step 9.

9. Enter and re-enter (confirm) the **Passphrase** used to encrypt the private key.

10. Click **OK**. The Device Certificates tab displays the imported certificate.

## Obtain a Certificate from an External CA

The advantage of obtaining a certificate from an external certificate authority (CA) is that the private key does not leave the firewall. To obtain a certificate from an external CA, generate a certificate signing request (CSR) and submit it to the CA. After the CA issues a certificate with the specified attributes, import it onto the firewall. The CA can be a well-known, public CA or an enterprise CA.

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of the certificate, Configure an OCSP Responder *before* generating the CSR.

| **Obtain a Certificate from an External CA** | | |
|---|---|---|
| Step 1 | Request the certificate from an external CA. | 1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.<br>In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.<br>2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in sub-step 6.<br>3. Click **Generate**.<br>4. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.<br>5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.<br>6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the device supports multiple virtual systems.<br>7. In the **Signed By** field, select **External Authority (CSR)**.<br>8. If applicable, select an **OCSP Responder**.<br>9. (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.<br> If you add a **Host Name** attribute, it is a best practice for it to match the **Common Name** (this is mandatory for GlobalProtect). The host name populates the Subject Alternative Name field of the certificate.<br>10. Click **Generate**. The Device Certificates tab displays the CSR with a Status of *pending*. |
| Step 2 | Submit the CSR to the CA. | 1. Select the CSR and click **Export** to save the .csr file to a local computer.<br>2. Upload the .csr file to the CA. |


© Palo Alto Networks, Inc.

| Obtain a Certificate from an External CA | |
|---|---|
| Step 3 Import the certificate. | 1. After the CA sends a signed certificate in response to the CSR, return to the **Device Certificates** tab and click **Import**.<br>2. Enter the **Certificate Name** used to generate the CSR in Step 1-4.<br>3. Enter the path and name of the PEM **Certificate File** that the CA sent, or **Browse** to it.<br>4. Click **OK**. The **Device Certificates** tab displays the certificate with a Status of *valid*. |
| Step 4 Configure the certificate. | 1. Click the certificate **Name**.<br>2. Select the check boxes that correspond to the intended use of the certificate on the firewall. For example, if the firewall will use this certificate to authenticate administrators who access the web interface, select the **Certificate for Secure Web GUI** check box.<br>3. Click **OK** and **Commit**. |

# Configure a Certificate Profile

Certificate profiles define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPSec VPN, Mobile Security Manager, and firewall/Panorama web interface access. The profiles specify which certificates to use, how to verify certificate revocation status, and how that status constrains access. Configure a certificate profile for each application.

It is a best practice to enable Online Certificate Status Protocol (OCSP) and/or Certificate Revocation List (CRL) status verification for certificate profiles. For details on these methods, see Certificate Revocation.

| Configure a Certificate Profile | | |
|---|---|---|
| Step 1 | Obtain the certificate authority (CA) certificates you will assign. | Perform one of the following steps to obtain the CA certificates you will assign to the profile. You must assign at least one.<br>• Generate a Certificate on the Firewall.<br>• Export a certificate from your enterprise CA and then import it onto the firewall (see Step 3). |
| Step 2 | Identify the certificate profile. | 1. In a firewall, select **Device > Certificate Management > Certificates Profile** and click **Add**.<br>In Panorama, select **Device > Certificate Management > Certificates Profile**, select a **Template** and click **Add**.<br>2. Enter a **Name** to identify the profile. The name is case-sensitive. It must be unique and use only letters, numbers, spaces, hyphens, and underscores. It can have up to 31 characters.<br>3. If the firewall supports multiple virtual systems, the dialog displays a **Location** drop-down. Select the virtual system where the profile will be available or select **Shared** to enable availability on all the virtual systems. |


© Palo Alto Networks, Inc.

| Configure a Certificate Profile | |
|---|---|
| Step 3 Assign one or more certificates. | Perform the following steps for each certificate:<br>1. In the CA Certificates table, click **Add**.<br>2. Select a **CA Certificate** from Step 1, or click **Import** and perform the following sub-steps:<br>a. Enter a **Certificate Name**.<br>b. Enter the path and name of the **Certificate File** you exported from your enterprise CA, or **Browse** to find the file.<br>c. Click **OK**.<br>3. Optionally, if the firewall uses OCSP to verify certificate revocation status, configure the following fields to override the default behavior. For most deployments, these fields do not apply.<br>• By default, the firewall uses the OCSP responder URL that you set in the procedure Configure an OCSP Responder. To override that setting, enter a **Default OCSP URL** (starting with *http://* or *https://*).<br>• By default, the firewall uses the certificate selected in the **CA Certificate** field to validate OCSP responses. To use a different certificate for validation, select it in the **OCSP Verify CA Certificate** field.<br>4. Click **OK**. The CA Certificates table displays the assigned certificate. |

| Configure a Certificate Profile | |
|---|---|
| Step 4 Define the methods for verifying certificate revocation status and the associated blocking behavior. | 1. Select **Use CRL** and/or **Use OCSP**. If you select both, the firewall first tries OCSP and falls back to the CRL method only if the OCSP responder is unavailable.<br>2. Depending on the verification method, enter the **CRL Receive Timeout** and/or **OCSP Receive Timeout**. These are the intervals (1-60 seconds) after which the firewall stops waiting for a response from the CRL/OCSP service.<br>3. Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies any session-blocking logic you define. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:<br>• If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.<br>• If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.<br>• If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value.<br>4. If you want the firewall to block sessions when the OCSP or CRL service returns a certificate revocation status of *unknown*, select the **Block session if certificate status is unknown** check box. Otherwise, the firewall proceeds with the session.<br>5. If you want the firewall to block sessions after it registers an OCSP or CRL request timeout, select the **Block session if certificate status cannot be retrieved within timeout** check box. Otherwise, the firewall proceeds with the session. |
| Step 5 Save and apply your entries. | Click **OK** and **Commit**. |


© Palo Alto Networks, Inc.

# Configure the Key Size for SSL Forward Proxy Server Certificates

When responding to a client in an SSL Forward Proxy session, the firewall creates a copy of the certificate that the destination server presents and uses the copy to establish a connection with the client. By default, the firewall generates certificates with the same key size as the certificate that the destination server presented. However, you can change the key size for the firewall-generated certificate as follows:

| Configure the Key Size for SSL Forward Proxy Server Certificates | |
|---|---|
| Step 1 | Select **Device > Setup > Session** and, in the Decryption Settings section, click **Forward Proxy Server Certificate Settings**. |
| Step 2 | Select a **Key Size**:<br>• **Defined by destination host**—The firewall determines the key size for the certificates it generates to establish SSL proxy sessions with clients based on the key size of the destination server certificate. If the destination server uses a 1,024-bit RSA key, the firewall generates a certificate with that key size and an SHA-1 hashing algorithm. If the destination server uses a key size larger than 1,024 bits (for example, 2,048 bits or 4,096 bits), the firewall generates a certificate that uses a 2,048-bit RSA key and SHA-256 algorithm. This is the default setting.<br>• **1024-bit RSA**—The firewall generates certificates that use a 1,024-bit RSA key and SHA-1 hashing algorithm regardless of the key size of the destination server certificates. As of December 31, 2013, public certificate authorities (CAs) and popular browsers have limited support for X.509 certificates that use keys of fewer than 2,048 bits. In the future, depending on security settings, when presented with such keys the browser might warn the user or block the SSL/TLS session entirely.<br>• **2048-bit RSA**—The firewall generates certificates that use a 2,048-bit RSA key and SHA-256 hashing algorithm regardless of the key size of the destination server certificates. Public CAs and popular browsers support 2,048-bit keys, which provide better security than the 1,024-bit keys.<br> Changing the key size setting clears the current certificate cache. |
| Step 3 | Click **OK** and **Commit**. |

# Revoke and Renew Certificates

▲ Revoke a Certificate
▲ Renew a Certificate

## Revoke a Certificate

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority (CA) that issued the certificate must revoke it. The following task describes how to revoke a certificate for which the firewall is the CA.

| **Revoke a Certificate** |
|---|
| 1. Select **Device > Certificate Management > Certificates > Device Certificates**. |
| 2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select the virtual system to which the certificate belongs. |
| 3. Select the certificate to revoke. |
| 4. Click **Revoke**. PAN-OS immediately sets the status of the certificate to revoked and adds the serial number to the Online Certificate Status Protocol (OCSP) responder cache or certificate revocation list (CRL). You need not perform a commit. |

## Renew a Certificate

If a certificate expires, or soon will, you can reset the validity period. If an external certificate authority (CA) signed the certificate and the firewall uses the Online Certificate Status Protocol (OCSP) to verify certificate revocation status, the firewall uses the OCSP responder information to update the certificate status (see Configure an OCSP Responder). If the firewall is the CA that issued the certificate, the firewall replaces it with a new certificate that has a different serial number but the same attributes as the old certificate.

| **Renew a Certificate** |
|---|
| 1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.<br>In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**. |
| 2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select the virtual system to which the certificate belongs. |
| 3. Select a certificate to renew and click **Renew**. |
| 4. Enter a **New Expiration Interval** (in days). |
| 5. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# Secure Keys with a Hardware Security Module

A hardware security module (HSM) is a physical device that manages digital keys. An HSM provides secure storage and generation of digital keys. It provides both logical and physical protection of these materials from non-authorized use and potential adversaries.

HSM clients integrated with Palo Alto Networks devices enable enhanced security for the private keys used in SSL/TLS decryption (both SSL forward proxy and SSL inbound inspection). In addition, you can use the HSM to encrypt device master keys.

The following topics describe how to integrate an HSM with your Palo Alto Networks devices:

- Set up Connectivity with an HSM
- Encrypt a Master Key Using an HSM
- Store Private Keys on an HSM
- Manage the HSM Deployment

# Set up Connectivity with an HSM

HSM clients are integrated with PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7050, and VM-Series firewalls and on Panorama (virtual appliance and M-100 appliance) for use with the following HSMs:

- SafeNet Network 5.2.1 or later
- Thales Nshield Connect 11.62 or later

The HSM server version must be compatible with these client versions. Refer to the HSM vendor documentation for the client-server version compatibility matrix.

The following topics describe how to set up connectivity between the firewall/Panorama and one of the supported HSMs:

▲ Set Up Connectivity with a SafeNet Network HSM

▲ Set Up Connectivity with a Thales Nshield Connect HSM

## Set Up Connectivity with a SafeNet Network HSM

To set up connectivity between the Palo Alto Networks device and a SafeNet Network HSM, you must specify the address of the HSM server and the password for connecting to it in the firewall configuration. In addition, you must register the firewall with the HSM server. Prior to beginning the configuration, make sure you have created a partition for the Palo Alto Networks devices on the HSM server.

HSM configuration is not synced between high availability firewall peers. Consequently, you must configure the HSM module separately on each of the peers.

In Active-Passive HA deployments, you must manually perform one failover to configure and authenticate each HA peer individually to the HSM. After this manual failover has been performed, user interaction is not required for the failover function.

| **Set up a Connectivity with a SafeNet Network HSM** | | |
|---|---|---|
| Step 1 | Configure the firewall to communicate with the SafeNet Network HSM. | 1. Log in to the firewall web interface and select **Device > Setup > HSM**.<br>2. Edit the Hardware Security Module Provider section and select **Safenet Luna SA** (SafeNet Network) as the **Provider Configured**.<br>3. Click **Add** and enter a **Module Name**. This can be any ASCII string up to 31 characters in length.<br>4. Enter the IPv4 address of the HSM module as the **Server Address**.<br>If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices.<br>5. (Optional) If configuring a high availability HSM configuration, select the **High Availability** check box and add the following: a value for **Auto Recovery Retry** and a **High Availability Group Name**.<br>If two HSM servers are configured, you should configure high availability. Otherwise the second HSM server is not used.<br>6. Click **OK** and **Commit.** |

 © Palo Alto Networks, Inc.

| Set up a Connectivity with a SafeNet Network HSM (Continued) | | |
|---|---|---|
| Step 2 | (Optional) Configure a service route to enable the firewall to connect to the HSM.<br><br>By default, the firewall uses the Management Interface to communicate with the HSM. To use a different interface, you must configure a service route. | 1. Select **Device > Setup > Services.**<br>2. Select **Service Route Configuration** from the Services Features area.<br>3. Select **Customize** from the Service Route Configuration area.<br>4. Select the **IPv4** tab.<br>5. Select **HSM** from the **Service** column.<br>6. Select an interface to use for HSM from the **Source Interface** drop-down.<br><br>If you select a dataplane connected port for HSM, issuing the `clear session all` CLI command will clear all existing HSM sessions, causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.<br>7. Click **OK** and **Commit.** |
| Step 3 | Configure the firewall to authenticate to the HSM. | 1. Select **Device > Setup > HSM**.<br>2. Select **Setup Hardware Security Module** in the Hardware Security Operations area.<br>3. Select the HSM **Server Name** from the drop-down.<br>4. Enter the **Administrator Password** to authenticate the firewall to the HSM.<br>5. Click **OK.**<br>The firewall attempts to perform an authentication with the HSM and displays a status message.<br>6. Click **OK.** |
| Step 4 | Register the firewall (the HSM client) with the HSM and assign it to a partition on the HSM.<br><br>If the HSM already has a firewall with the same `<cl-name>` registered, you must remove the duplicate registration using the following command before registration will succeed:<br><br>`client delete -client <cl-name>`<br><br>where `<cl-name>` is the name of the client (firewall) registration you want to delete. | 1. Log in to the HSM from a remote system.<br>2. Register the firewall using the following command:<br>`client register -c <cl-name> -ip <fw-ip-addr>`<br>where `<cl-name>` is a name that you assign to the firewall for use on the HSM and `<fw-ip-addr>` is the IP address of the firewall that is being configured as an HSM client.<br>3. Assign a partition to the firewall using the following command:<br>`client assignpartition -c <cl-name> -p <partition-name>`<br>where `<cl-name>` is the name assigned to the firewall in the `client register` command and `<partition-name>` is the name of a previously configured partition that you want to assign to the firewall. |

| Set up a Connectivity with a SafeNet Network HSM (Continued) | | |
|---|---|---|
| Step 5 | Configure the firewall to connect to the HSM partition. | 1. Select **Device > Setup > HSM**.<br>2. Click the **Refresh** icon.<br>3. Select the **Setup HSM Partition** in the Hardware Security Operations area.<br>4. Enter the **Partition Password** to authenticate the firewall to the partition on the HSM.<br>5. Click **OK**. |
| Step 6 | (Optional) Configure an additional HSM for high availability (HA). | 1. Follow Step 1 through Step 5 to add an additional HSM for high availability (HA).<br>This process adds a new HSM to the existing HA group.<br>2. If you remove an HSM from your configuration, repeat Step 5.<br>This will remove the deleted HSM from the HA group. |
| Step 7 | Verify connectivity with the HSM. | 1. Select **Device > Setup > HSM**.<br>2. Check the **Status** of the HSM connection:<br>Green—HSM is authenticated and connected.<br>Red—HSM was not authenticated or network connectivity to the HSM is down.<br>3. View the following columns in Hardware Security Module Status area to determine authentication status:<br>**Serial Number**—The serial number of the HSM partition if the HSM was successfully authenticated.<br>**Partition—**The partition name on the HSM that was assigned on the firewall.<br>**Module State**—The current operating state of the HSM. It always has the value Authenticated if the HSM is displayed in this table. |

## Set Up Connectivity with a Thales Nshield Connect HSM

The following workflow describes how to configure the firewall to communicate with a Thales Nshield Connect HSM. This configuration requires that you set up a remote filesystem (RFS) to use as a *hub* to sync key data for all firewalls in your organization that are using the HSM.

HSM configuration is not synced between high availability firewall peers. Consequently, you must configure the HSM module separately on each of the peers.

If the high availability firewall configuration is in Active-Passive mode, you must manually perform one failover to configure and authenticate each HA peer individually to the HSM. After this manual failover has been performed, user interaction is not required for the failover function.

 © Palo Alto Networks, Inc.

| Set up Connectivity with a Thales Nshield Connect HSM | | |
|---|---|---|
| Step 1 | Configure the Thales Nshield Connect server as the firewall's HSM provider. | 1. From the firewall web interface, select **Device > Setup > HSM** and edit the Hardware Security Module Provider section.<br>2. Select **Thales Nshield Connect** as the **Provider Configured**.<br>3. Click **Add** and enter a **Module Name**. This can be any ASCII string up to 31 characters in length.<br>4. Enter the IPv4 address as the **Server Address** of the HSM module.<br>If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices.<br>5. Enter the IPv4 address of the **Remote Filesystem Address**.<br>6. Click **OK** and **Commit**. |
| Step 2 | (Optional) Configure a service route to enable the firewall to connect to the HSM.<br>By default, the firewall uses the Management Interface to communicate with the HSM. To use a different interface, you must configure a service route. | 1. Select **Device > Setup > Services.**<br>2. Select **Service Route Configuration** from the Services Features area.<br>3. Select **Customize** from the Service Route Configuration area.<br>4. Select the **IPv4** tab.<br>5. Select **HSM** from the **Service** column.<br>6. Select an interface to use for HSM from the **Source Interface** drop-down.<br>If you select a dataplane connected port for HSM, issuing the `clear session all` CLI command will clear all existing HSM sessions, causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.<br>7. Click **OK** and **Commit.** |
| Step 3 | Register the firewall (the HSM client) with the HSM server.<br>This step briefly describes the procedure for using the front panel interface of the Thales Nshield Connect HSM. For more details, consult the Thales documentation. | 1. Log in to the front panel display of the Thales Nshield Connect HSM unit.<br>2. On the unit front panel, use the right-hand navigation button to select **System > System configuration > Client config > New client**.<br>3. Enter the IP address of the firewall.<br>4. Select **System > System configuration > Client config > Remote file system** and enter the IP address of the client computer where you set up the remote file system. |

| Set up Connectivity with a Thales Nshield Connect HSM (Continued) | | |
|---|---|---|
| Step 4 | Set up the remote filesystem to accept connections from the firewall. | 1. Log in to the remote filesystem (RFS) from a Linux client.<br>2. Obtain the electronic serial number (ESN) and the hash of the $K_{NETI}$ key. The $K_{NETI}$ key authenticates the module to clients:<br>`anonkneti <ip-address>`<br>where `<ip-address>` is the IP address of the HSM.<br>The following is an example:<br>`anonkneti 192.0.2.1`<br>`B1E2-2D4C-E6A2 5a2e5107e70d525615a903f6391ad72b1c03352c`<br>In this example, **`B1E2-2D4C-E6A2`** is the ESM and 5a2e5107e70d525615a903f6391ad72b1c03352c is the hash of the $K_{NETI}$ key.<br>3. Use the following command from a superuser account to perform the remote filesystem setup:<br>`rfs-setup --force <ip-address> <ESN> <hash-Kneti-key>`<br>where `<ip-address>` is the IP address of the HSM,<br>`<ESN>` is the electronic serial number (ESN) and<br>`<hash-Kneti-key>` is the hash of the KNETI key.<br>The following example uses the values obtained in this procedure:<br>`rfs-setup --force <192.0.2.1> <B1E2-2D4C-E6A2>`<br>`<5a2e5107e70d525615a903f6391ad72b1c03352c>`<br>4. Use the following command to permit client submit on the Remote Filesystem:<br>`rfs-setup --gang-client --write-noauth <FW-IPaddress>`<br>where `<FW-IPaddress>` is the IP address of the firewall. |
| Step 5 | Configure the firewall to authenticate to the HSM. | 1. From the firewall web interface, select **Device > Setup > HSM**.<br>2. Select **Setup Hardware Security Module** in the Hardware Security Operations area.<br>3. Click **OK.**<br>The firewall attempts to perform an authentication with the HSM and displays a status message.<br>4. Click **OK.** |
| Step 6 | Synchronize the firewall with the remote filesystem. | 1. Select the **Device > Setup > HSM**.<br>2. Select **Synchronize with Remote Filesystem** in the Hardware Security Operations section. |


© Palo Alto Networks, Inc.

| Set up Connectivity with a Thales Nshield Connect HSM (Continued) | | |
|---|---|---|
| Step 7 | Verify that the firewall can connect to the HSM. | 1. Select **Device > Setup > HSM**.<br>2. Check the Status indicator to verify that the firewall is connected to the HSM:<br>Green—HSM is authenticated and connected.<br>Red—HSM was not authenticated or network connectivity to the HSM is down.<br>3. View the following columns in Hardware Security Module Status section to determine authentication status.<br>**Name**: The name of the HSM attempting to be authenticated.<br>**IP address:** The IP address of the HSM that was assigned on the firewall.<br>**Module State**: The current operating state of the HSM: **Authenticated** or **Not Authenticated**. |

# Encrypt a Master Key Using an HSM

A master key is configured on a Palo Alto Networks firewall to encrypt all private keys and passwords. If you have security requirements to store your private keys in a secure location, you can encrypt the master key using an encryption key that is stored on an HSM. The firewall then requests the HSM to decrypt the master key whenever it is required to decrypt a password or private key on the firewall. Typically, the HSM is located in a highly secure location that is separate from the firewall for greater security.

The HSM encrypts the master key using a wrapping key. To maintain security, this encryption key must occasionally be changed. For this reason, a command is provided on the firewall to rotate the wrapping key which changes the master key encryption. The frequency of this wrapping key rotation depends on your application.

Master key encryption using an HSM is not supported on firewalls configured in FIPS or CC mode.

The following topics describe how to encrypt the master key initially and how to refresh the master key encryption:

- Encrypt the Master Key
- Refresh the Master Key Encryption

## Encrypt the Master Key

If you have not previously encrypted the master key on a device, use the following procedure to encrypt it. Use this procedure for first time encryption of a key, or if you define a new master key and you want to encrypt it. If you want to refresh the encryption on a previously encrypted key, see Refresh the Master Key Encryption.

| Encrypt a Master Key Using an HSM | |
|---|---|
| Step 1 | Select **Device > Master Key and Diagnostics.** |
| Step 2 | Specify the key that is currently used to encrypt all of the private keys and passwords on the firewall in the **Master Key** field. |
| Step 3 | If changing the master key, enter the new master key and confirm. |
| Step 4 | Select the **HSM** check box.<br>**Life Time:** The number of days and hours after which the master key expires (range 1-730 days).<br>**Time for Reminder:** The number of days and hours before expiration when the user is notified of the impending expiration (range 1-365 days). |
| Step 5 | Click **OK**. |

## Refresh the Master Key Encryption

As a best practice, refresh the master key encryption on a regular basis by rotating the master key wrapping key on the HSM. This command is the same for both the SafeNet Network and Thales Nshield Connect HSMs.

 © Palo Alto Networks, Inc.

**Refresh the Master Key Encryption**

1. Use the following CLI command to rotate the wrapping key for the master key on an HSM:

   ```
   > request hsm mkey-wrapping-key-rotation
   ```

   If the master key is encrypted on the HSM, the CLI command will generate a new wrapping key on the HSM and encrypt the master key with the new wrapping key.

   If the master key is not encrypted on the HSM, the CLI command will generate new wrapping key on the HSM for future use.

   The old wrapping key is not deleted by this command.

## Store Private Keys on an HSM

For added security, the private keys used to enable SSL/TLS decryption—both SSL forward proxy and SSL inbound inspection—can be secured with an HSM as follows:

- **SSL forward proxy**—The private key in the CA certificate that is used to sign certificates in SSL/TLS forward proxy operations can be stored on the HSM. The firewall will then send the certificates it generates during SSL/TLS forward proxy operations to the HSM for signing before forwarding them on to the client.
- **SSL inbound inspection**—The private keys for the internal servers for which you are doing SSL/TLS inbound inspection can be stored on the HSM.

For instructions on importing the private keys onto the HSM, refer to the documentation from your HSM provider. After the required keys are on the HSM, you can configure the firewall to locate the keys as follows:

| **Store Private Keys on an HSM** | | |
|---|---|---|
| Step 1 | Import the private keys used in your SSL forward proxy and/or SSL inbound inspection deployments onto the HSM. | For instructions on importing the private keys onto the HSM, refer to the documentation from your HSM provider. |
| Step 2 | (Thales Nshield Connect only) Sync the key data from the HSM remote file system to the firewall. | 1. From the firewall web interface, select **Device > Setup > HSM**.<br>2. Select **Synchronize with Remote Filesystem** in the Hardware Security Operations section. |
| Step 3 | Import the certificate(s) that correspond to the private key(s) you are storing on the HSM onto the firewall. | 1. From the firewall web interface, select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Click **Import**.<br>3. Enter the **Certificate Name**.<br>4. Enter filename of the **Certificate File** you imported to the HSM.<br>5. Select the appropriate **File Format** from the drop-down.<br>6. Select the **Private Key resides on Hardware Security Module** check box.<br>7. Click **OK** and **Commit**. |
| Step 4 | (Forward trust certificates only) Enable the certificate for use in SSL/TLS Forward Proxy. | 1. Select the **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Locate the certificate you imported in Step 3.<br>3. Select the **Forward Trust Certificate** check box.<br>4. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

| Store Private Keys on an HSM (Continued) | | |
|---|---|---|
| Step 5 | Verify that the certificate has been successfully imported to the firewall. | 1. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Locate the certificate you imported in Step 3.<br>3. In the Key column, notice the following:<br>If a Lock icon is displayed, the private key for the certificate can be found on the HSM.<br>If an Error icon is displayed, the private key is not imported to the HSM or the HSM is not properly authenticated or connected. |

## Manage the HSM Deployment

| Manage HSM | |
|---|---|
| • View the HSM configuration settings. | Select **Device > Setup > HSM**. |
| • Display detailed HSM information. | Select **Show Detailed Information** from the Hardware Security Operations section.<br>Information regarding the HSM servers, HSM HA status, and HSM hardware is displayed. |
| • Export Support file. | Select **Export Support File** from the Hardware Security Operations section.<br>A test file is created to help customer support when addressing a problem with an HSM configuration on the firewall. |
| • Reset HSM configuration. | Select **Reset HSM Configuration** from the Hardware Security Operations section.<br>Selecting this option removes all HSM connections. All authentication procedures must be repeated after using this option. |

 © Palo Alto Networks, Inc.

# High Availability

High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in HP pairs provides redundancy and allows you to ensure business continuity.

The Palo Alto Networks firewalls support stateful active/passive or active/active high availability with session and configuration synchronization. Some models of the firewall, such as the VM-Series firewall and the PA-200, only support HA lite without session synchronization capability. The following topics provide more information about high availability and how to configure it in your environment.

- HA Overview
- HA Concepts
- Set Up Active/Passive HA
- HA Resources

# HA Overview

On Palo Alto Networks firewalls, you can set up two firewalls as an HA pair. HA allows you to minimize downtime by making sure that an alternate firewall is available in the event that the peer firewall fails. The firewalls in an HA pair use dedicated or in-band HA ports on the firewall to synchronize data—network, object, and policy configurations—and to maintain state information. Firewall-specific configuration such as management interface IP address or administrator profiles, HA specific configuration, log data, and the Application Command Center (ACC) information is not shared between peers. For a consolidated application and log view across the HA pair, you must use Panorama, the Palo Alto Networks centralized management system.

When a failure occurs on one firewall in the pair and the other firewall takes over the task of securing traffic, the event is called a *failover*. The conditions that trigger a failover are:

- One or more of the monitored interfaces fail. (Link Monitoring)
- One or more of the destinations specified on the firewall cannot be reached. (Path Monitoring)
- The firewall does not respond to heartbeat polls. (Heartbeat Polling and Hello messages)

After you understand the HA Concepts, continue to Set Up Active/Passive HA.

 © Palo Alto Networks, Inc.

# HA Concepts

The following topics provide conceptual information about how HA works on a Palo Alto Networks firewall:


- HA Modes
- HA Links and Backup Links
- Device Priority and Preemption
- Failover Triggers
- HA Timers


## HA Modes

You can set up the firewalls for HA in two modes:

- **Active/Passive**— One firewall actively manages traffic while the other is synchronized and ready to transition to the active state, should a failure occur. In this configuration, both firewalls share the same configuration settings, and one actively manages traffic until a path, link, system, or network failure occurs. When the active firewall fails, the passive firewall takes over seamlessly and enforces the same policies to maintain network security. Active/passive HA is supported in the virtual wire, Layer 2 and Layer 3 deployments. For information on setting up your firewall in an active/passive configuration, see Configure Active/Passive HA.

The PA-200 and the VM-Series firewalls support a lite version of active/passive HA. HA lite provides configuration synchronization and some runtime data synchronization such as IPSec security associations. It does not support any session synchronization, and therefore, HA Lite does not offer stateful failover.

- **Active/Active**— Both the firewalls in the pair are active and processing traffic, and work synchronously to handle session setup and session ownership. The active/active deployment is supported in virtual wire and Layer 3 deployments, and is only recommended for networks with asymmetric routing. For information on setting up the firewalls in an active/active configuration, refer to the Active/Active High Availability Tech Note.

## HA Links and Backup Links

The firewalls in an HA pair use HA links to synchronize data and maintain state information. Some models of the firewall have dedicated HA ports—Control link (HA1) and Data link (HA2), while others require you to use the in-band ports as HA links.

On firewalls with dedicated HA ports such as the PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7050 firewalls (see HA Ports on the PA-7050 Firewall), use the dedicated HA ports to manage communication and synchronization between the firewalls. For firewalls without dedicated HA ports such as the PA-200, PA-500, and PA-2000 Series firewalls, as a best practice use the management port for the HA1 link to allow for a direct connection between the management planes on the firewalls, and an in-band port for the HA2 link.

The HA1 and HA2 links provide synchronization for functions that reside on the management plane. Using the dedicated HA interfaces on the management plane is more efficient than using the in-band ports as this eliminates the need to pass the synchronization packets over the dataplane.

- **Control Link**: The HA1 link is used to exchange hellos, heartbeats, and HA state information, and management plane sync for routing, and User-ID information. This link is also used to synchronize configuration changes on either the active or passive firewall with its peer. The HA1 link is a Layer 3 link and requires an IP address.

  **Ports used for HA1**: TCP port 28769 and 28260 for clear text communication; port 28 for encrypted communication (SSH over TCP).

- **Data Link**: The HA2 link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between firewalls in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active firewall to the passive firewall. The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.

  **Ports used for HA2:** The HA data link can be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets.

  Additionally, an HA3 link is used in Active/Active HA deployments. When there is an asymmetric route, the HA3 link is used for forwarding packets to the HA peer that owns the session. The HA3 link is a Layer 2 link and it does not support Layer 3 addressing or encryption.

- **Backup Links**: Provide redundancy for the HA1 and the HA2 links. In-band ports are used as backup links for both HA1 and HA2. Consider the following guidelines when configuring backup HA links:
  - The IP addresses of the primary and backup HA links must not overlap each other.
  - HA backup links must be on a different subnet from the primary HA links.
  - HA1-backup and HA2-backup ports must be configured on separate physical ports. The HA1-backup link uses port 28770 and 28260.

Palo Alto Networks recommends enabling heartbeat backup (uses port 28771 on the MGT interface) if you use an in-band port for the HA1 or the HA1 backup links.

- **Packet-Forwarding Link**: In addition to the HA1 and HA2 links, an active/active deployment also requires a dedicated HA3 link. The firewalls use this link for forwarding packets to the peer during session setup and asymmetric traffic flow. The HA3 link is a Layer 2 link that uses MAC-in-MAC encapsulation; it does not support Layer 3 addressing or encryption. You can configure aggregate interfaces on the PA-3000 Series, PA-4000 Series, PA-5000 Series and PA-7000 Series firewalls as an HA3 link. The aggregate interfaces can also provide redundancy for the HA3 link; you cannot configure backup links for the HA3 link.

## HA Ports on the PA-7050 Firewall

For HA connectivity on the PA-7050, use the following table for details on which ports on the Switch Management Card (SMC) are mandated and where ports on the Network Processing Card (NPC) are suitable.

 © Palo Alto Networks, Inc.

| HA Links and Backup Links | Ports on the SMC | Description |
|---|---|---|
| **Control Link** | HA1-A<br>Speed: Ethernet 10/100/1000 | Used for HA control and synchronization. Connect this port directly from the HA1-A port on the first firewall to the HA1-A on the second firewall in the pair, or connect them together through a switch or router.<br>HA1 cannot be configured on NPC data ports or the MGT port. |
| **Control Link Backup** | HA1-B<br>Speed: Ethernet 10/100/1000 port | Used for HA control and synchronization as a backup for HA1-A. Connect this port directly from the HA1-B port on the first firewall to the HA1-B on the second firewall in the pair, or connect them together through a switch or router.<br>HA1 Backup cannot be configured on NPC data ports or the MGT port. |
| **Data Link**<br><br>**Data Link Backup** | HSCI-A<br><br>HSCI-B | The High Speed Chassis Interconnect (HSCI) ports are Quad Port SFP (QSFP) interfaces which are used to connect two PA-7000 Series firewalls in an HA configuration. Each port is comprised of four 10 gigabit links internally for a combined speed of 40 gigabits.<br>The HSCI ports are not routable and must be connected directly to each other. The HSCI-A on the first chassis connects directly to HSCI-A on the second chassis and HSCI-B on the first chassis connects to HSCI-B on the second chassis. This will provide full 80 gigabit transfer rates. In software, both ports (HSCI-A and HSCI-B) are treated as one HA interface.<br>Palo Alto Networks recommends using the dedicated HSCI ports for the HA2 link; the HA3 link must use the HSCI port. If the firewalls are deployed in:<br>• an active/active configuration, the HA3 link must use the HSCI port. The HA2 link and HA2 backup links can use the HSCI port or data ports on the NPC.<br>• an active/passive configuration, you can configure a data port on the NPC for the HA2 link or the HA2 backup link, if needed. |

For an overview of the Modules and Interface cards on the PA-7050 firewall, refer to the PA-7050 Hardware Reference Guide.

## Device Priority and Preemption

The firewalls in an HA pair can be assigned a *device priority* value to indicate a preference for which firewall should assume the active role and manage traffic. If you need to use a specific firewall in the HA pair for actively securing traffic, you must enable the preemptive behavior on both the firewalls and assign a device priority value for each firewall. The firewall with the lower numerical value, and therefore *higher priority*, is designated as active and manages all traffic on the network. The other firewall is in a passive state, and synchronizes configuration and state information with the active firewall so that it is ready to transition to an active state should a failure occur.

By default, preemption is disabled on the firewalls and must be enabled on both firewalls. When enabled, the preemptive behavior allows the firewall with the *higher priority* (lower numerical value) to resume as active after it recovers from a failure. When preemption occurs, the event is logged in the system logs.

## Failover Triggers

When a failure occurs on one firewall and the peer takes over the task of securing traffic, the event is called a *failover*. A failover is triggered when a monitored metric on a firewall in an HA pair fails. The metrics that are monitored for detecting a device failure are:

- **Heartbeat Polling and Hello messages**

  The firewalls use hello message and heartbeats to verify that the peer firewall is responsive and operational. Hello messages are sent from one peer to the other at the configured *Hello Interval* to verify the state of the firewall. The heartbeat is an ICMP ping to the HA peer over the control link, and the peer responds to the ping to establish that the firewalls are connected and responsive. By default, the interval for the heartbeat is 1000 milliseconds. For details on the HA timers that trigger a failover, see HA Timers.

- **Link Monitoring**

  The physical interfaces to be monitored are grouped into a link group and their state (link up or link down) is monitored. A link group can contain one or more physical interfaces. A firewall failure is triggered when any or all of the interfaces in the group fail. The default behavior is failure of any one link in the link group will cause the firewall to change the HA state to non-functional to indicate a failure of a monitored object.

- **Path Monitoring**

  Monitors the full path through the network to mission-critical IP addresses. ICMP pings are used to verify reachability of the IP address. The default interval for pings is 200ms. An IP address is considered unreachable when 10 consecutive pings (the default value) fail, and a firewall failure is triggered when any or all of the IP addresses monitored become unreachable. The default behavior is any one of the IP addresses becoming unreachable will cause the firewall to change the HA state to non-functional to indicate a failure of a monitored object.

In addition to the failover triggers listed above, a failover also occurs when the administrator places the firewall is a suspended state or if preemption occurs.

On the PA-3000 Series, PA-5000 Series, and PA-7050 firewalls, a failover can occur when an internal health check fails. This health check is not configurable and is enabled to verify the operational status for all the components within the firewall.

## HA Timers

High Availability (HA) timers are used to detect a firewall failure and trigger a failover. To reduce the complexity in configuring HA timers, you can select from three profiles have been added: **Recommended**, **Aggressive** and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.

 © Palo Alto Networks, Inc.

Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

The following table describes each timer included in the profiles and the current preset values across the different hardware models; these values are for current reference only and can change in a subsequent release.

**Recommended/Aggressive HA Timer Values by Platform**

| Timers | Description | PA-7050<br>PA-5000 Series<br>PA-4000 Series<br>PA-3000 Series | PA-2000 Series<br>PA-500 Series<br>PA-200 Series<br>VM-Series | Panorama VM<br>M-100 |
|---|---|---|---|---|
| Monitor fail hold up time | Interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is recommended to avoid an HA failover due to the occasional flapping of neighboring devices. | 0/0 | 0/0 | 0/0 |
| Preemption hold time | Time a passive or active-secondary firewall will wait before taking over as the active or active-primary firewall. | 1/1 | 1/1 | 1/1 |
| Heartbeat interval | Frequency at which the HA peers exchange heartbeat messages in the form of an ICMP ping. | 1000/1000 | 2000/1000 | 2000/1000 |
| Promotion hold time | Time that the passive firewall (in active/passive mode) or the active-secondary firewall (in active/active mode) will wait before taking over as the active or active-primary firewall after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made. | 2000/500 | 2000/500 | 2000/500 |

| Timers | Description | PA-7050<br>PA-5000 Series<br>PA-4000 Series<br>PA-3000 Series | PA-2000 Series<br>PA-500 Series<br>PA-200 Series<br>VM-Series | Panorama VM<br>M-100 |
|---|---|---|---|---|
| Additional master hold up time | This time interval is applied to the same event as Monitor Fail Hold Up Time (range 0-60000 ms, default 500 ms). The additional time interval is applied only to the active firewall in active/passive mode and to the active-primary firewall in active/active mode. This timer is recommended to avoid a failover when both firewalls experience the same link/path monitor failure simultaneously. | 500/500 | 500/500 | 7000/5000 |
| Hello interval | Interval in milliseconds between the hello packets that are sent to verify that the HA functionality on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms. | 8000/8000 | 8000/8000 | 8000/8000 |
| Maximum no. of flaps | A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. This value indicates the maximum number of flaps that are permitted before the firewall is determined to be suspended and the passive firewall takes over (range 0-16, default 3). | 3/3 | 3/3 | Not Applicable |

 © Palo Alto Networks, Inc.

# Set Up Active/Passive HA


- Prerequisites for Active/Passive HA
- Configuration Guidelines for Active/Passive HA
- Configure Active/Passive HA
- Define HA Failover Conditions
- Verify Failover

## Prerequisites for Active/Passive HA

To set up high availability on your Palo Alto Networks firewalls, you need a pair of firewalls that meet the following requirements:

- **The same model**—Both the firewalls in the pair must be of the same hardware model or virtual machine model.
- **The same PAN-OS version**—Both the firewalls must be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases. They must also both have the same multiple virtual systems capability (single or multi vsys).
- **The same type of interfaces**—Dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
  - Determine the IP address for the HA1 (control) connection between the HA peers. The HA1 IP address for both peers must be on the same subnet if they are directly connected or are connected to the same switch.

    For firewall without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both firewalls. However, because the management ports will not be directly cabled between the firewalls, make sure that you have a route that connects these two interfaces across your network.
  - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use Layer 3 *only* if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
- **The same set of licenses**—Licenses are unique to each firewall and cannot be shared between the firewalls. Therefore, you must license both firewalls identically. If both firewalls do not have an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.

If you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration, it is recommended that you Reset the Firewall to Factory Default Settings on the new firewall. This will ensure that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary firewall to the newly introduced firewall with the clean config.

 © Palo Alto Networks, Inc.

## Configuration Guidelines for Active/Passive HA

To set up an active (PeerA) passive (PeerB) pair in HA, you must configure some options identically on both firewalls and some independently (non-matching) on each firewall. These HA settings are not synchronized between the firewalls. For details on what is/is not synchronized, refer to HA Synchronization.

To proceed with the instructions on configuring the firewalls in HA, see Configure Active/Passive HA.

The following table lists the settings that you must configure identically on both firewalls:

| Identical Configuration Settings on PeerA and PeerB |
|---|
| • HA must be enabled on both firewalls.<br>• Both firewall must have the same Group ID value. The Group ID value is used to create a virtual MAC address for all the configured interfaces. The format of the virtual MAC is 00-1B-17:00: xx: yy where<br>00-1B-17: vendor ID; 00: fixed; xx: HA group ID; yy: interface ID.<br>When a new active firewall takes over, Gratuitous ARPs are sent from each of the connected interfaces of the new active member to inform the connected Layer 2 switches of the virtual MAC address' new location.<br>• If using in-band ports, the interfaces for the HA1 and HA2 links must be set to type HA.<br>• The HA mode must be set to Active Passive.<br>• If required, preemption must be enabled on both firewalls. The device priority value, however, must not be identical.<br>• If required, encryption on the HA1 link (for communication between the HA peers) must be configured on both firewalls.<br>• Based on the combination of HA1 and HA1 Backup ports you are using, use the following recommendations to decide whether you should enable heartbeat backup:<br>• HA1: Dedicated HA1 port<br>HA1 Backup: In-band port<br>Recommendation: Enable Heartbeat Backup<br>• HA1: Dedicated HA1 port<br>HA1 Backup: Management port<br>Recommendation: Do not enable Heartbeat Backup<br>• HA1: In-band port<br>HA1 Backup: In-band port<br>Recommendation: Enable Heartbeat Backup<br>• HA1: Management port<br>HA1 Backup: In-band port<br>Recommendation: Do not enable Heartbeat Backup |

The following table lists the settings that must be configured independently on each firewall:

| Independent Configuration Settings | PeerA | PeerB |
|---|---|---|
| Control Link | IP address of the HA1 link configured on this firewall (PeerA). | IP address of the HA1 link configured on this firewall (PeerB). |
| | For firewalls without dedicated HA ports, use the management port IP address for the control link. | |
| Data Link<br>The data link information is synchronized between the firewalls after HA is enabled and the control link is established between the firewalls. | By default, the HA2 link uses Ethernet/Layer 2.<br>If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerA). | By default, the HA2 link uses Ethernet/Layer 2.<br>If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerB). |
| Device Priority (required, if preemption is enabled) | The firewall you plan to make active must have a lower numerical value than its peer. So, if Peer A is to function as the active firewall, keep the default value of 100 and increment the value on PeerB. | If PeerB is passive, set the device priority value to a number larger than that on PeerA. For example, set the value to 110. |
| Link Monitoring—Monitor one or more physical interfaces that handle vital traffic on this firewall and define the failure condition. | Select the physical interfaces on the firewall that you would like to monitor and define the failure condition (all or any) to trigger a failover. | Pick a similar set of physical interfaces that you would like to monitor on this firewall and define the failure condition (all or any) to trigger a failover. |
| Path Monitoring—Monitor one or more destination IP addresses that the firewall can use ICMP pings to ascertain responsiveness. | Define the failure condition (all or any), ping interval and the ping count. This is particularly useful for monitoring the availability of other interconnected networking devices. For example, monitor the availability of a router that connects to a server, connectivity to the server itself, or some other vital device that is in the flow of traffic.<br>Make sure that the node/device that you are monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a a path monitoring failure and trigger a failover. | Pick a similar set of devices or destination IP addresses that can be monitored for determining the failover trigger for PeerB. Define the failure condition (all or any), ping interval and the ping count. |

 © Palo Alto Networks, Inc.

## Configure Active/Passive HA

The following procedure shows how to configure a pair of firewalls in an active/passive deployment as depicted in the following example topology.

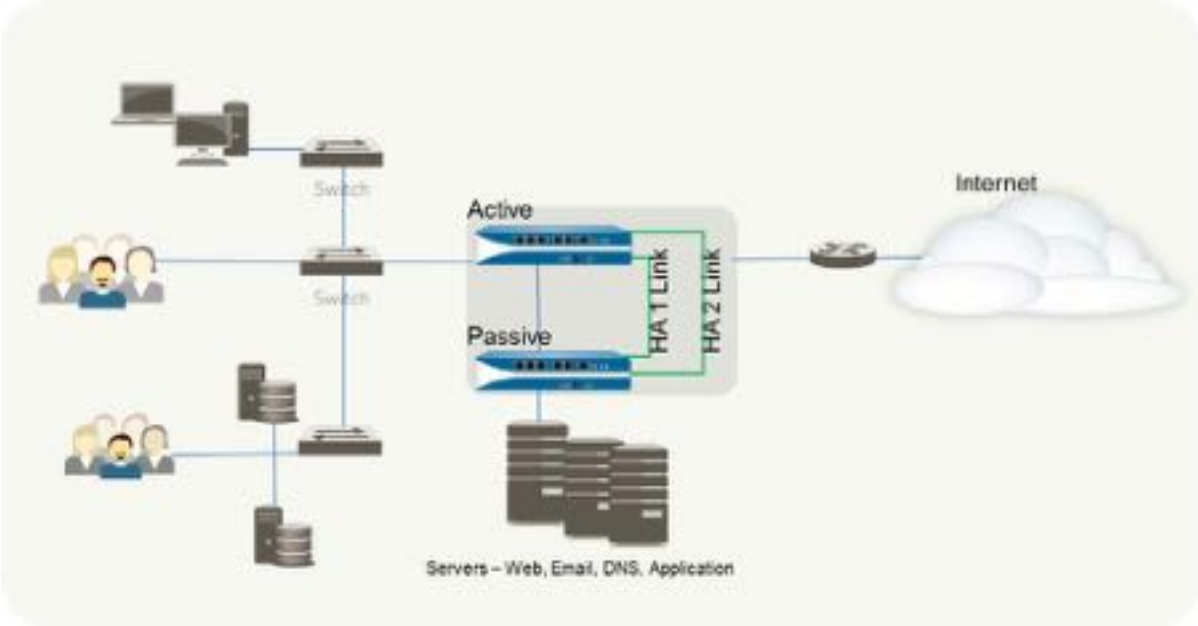


| Connect and Configure the Firewalls | |
|---|---|
| Step 1 Connect the HA ports to set up a physical connection between the firewalls. | • For firewalls with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on the HA firewalls. Use a crossover cable if the firewalls are directly connected to each other.<br>• For firewalls without dedicated HA ports, select two data interfaces for the HA2 link and the backup HA1 link. Then, use an Ethernet cable to connect these in-band HA interfaces across both firewalls.<br>Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network. |
| Pick a firewall in the pair and complete the following steps: | |
| Step 2 Enable ping on the management port.<br>Enabling ping allows the management port to exchange heartbeat backup information. | 1. Select **Device > Setup > Management** and edit the Management Interface Settings.<br>2. Select **Ping** as a service that is permitted on the interface. |

| Connect and Configure the Firewalls (Continued) | | |
|---|---|---|
| Step 3 | If the firewall does not have dedicated HA ports, set up the data ports to function as HA ports.<br><br>For firewalls with dedicated HA ports continue to the next step. | 1. Select **Network > Interfaces.**<br>2. Confirm that the link is up on the ports that you want to use.<br>3. Select the interface and set the interface type to HA.<br> 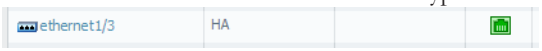 <br>4. Set the **Link Speed** and **Link Duplex** settings, as appropriate. |
| Step 4 | Set the HA mode and group ID. | 1. Select **Device > High Availability > General** and edit the Setup section.<br>2. Set a **Group ID** and optionally a **Description** for the pair. The Group ID uniquely identifies each HA pair on your network. If you have have multiple HA pairs that share the same broadcast domain you must set a unique Group ID for each pair.<br>3. Set the mode to **Active Passive**. |
| Step 5 | Set up the control link connection.<br><br>This example shows an in-band port that is set to interface type HA.<br><br>For firewalls that use the management port as the control link, the IP address information is automatically pre-populated. | 1. In **Device > High Availability > General**, edit the Control Link (HA1) section.<br>2. Select the Port that you have cabled for use as the HA1 link.<br>3. Set the IPv4/IPv6 Address and Netmask.<br>4. If the HA1 interfaces are on separate subnets, enter the IP address of the Gateway. Do not add a gateway address if the firewalls are directly connected.<br> |
| Step 6 | (Optional) Enable encryption for the control link connection.<br><br>This is typically used to secure the link if the two firewalls are not directly connected, that is if the ports are connected to a switch or a router. | 1. Export the HA key from a firewall and import it into the peer firewall.<br>a. Select **Device > Certificate Management > Certificates**.<br>b. Select **Export HA key**. Save the HA key to a network location that the peer can access.<br>c. On the peer, select **Device > Certificate Management > Certificates,** and select **Import HA key** to browse to the location that you saved the key and import it in to the peer.<br>2. Select **Device > High Availability > General**, edit the Control Link (HA1) section.<br>3. Select **Encryption Enabled**. |

 © Palo Alto Networks, Inc.

| Connect and Configure the Firewalls (Continued) | |
|---|---|
| Step 7 Set up the backup control link connection. | 1. In **Device > High Availability > General**, edit the Control Link (HA1 Backup) section.<br>2. Select the HA1 backup interface and set the **IPv4/IPv6 Address** and **Netmask**.<br> |
| Step 8 Set up the data link connection (HA2) and the backup HA2 connection between the firewalls. | 1. In **Device > High Availability > General**, edit the Data Link (HA2) section.<br>2. Select the **Port** to use for the data link connection.<br>3. Select the **Transport** method. The default is **ethernet**, and will work when the HA pair is connected directly or through a switch. If you need to route the data link traffic through the network, select **IP** or **UDP** as the transport mode.<br>4. If you use IP or UDP as the transport method, enter the **IPv4/IPv6 Address** and **Netmask**.<br> 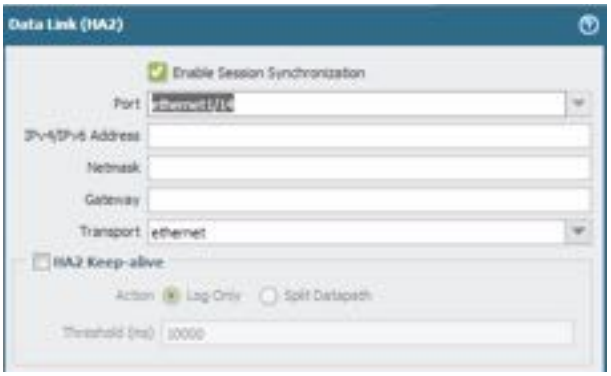 <br>5. Verify that **Enable Session Synchronization** is selected.<br>6. Select **HA2 Keep-alive** to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set (default is 10000 ms), the defined action will occur. For active/passive configuration, a critical system log message is generated when an HA2 keep-alive failure occurs.<br>You can configure the HA2 keep-alive option on both firewalls, or just one firewall in the HA pair. If you only enable the option on one peer, only that firewall will send the keep-alive messages. The other firewall will be notified if a failure occurs.<br>7. Edit the Data Link (HA2 Backup) section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**. |

| Connect and Configure the Firewalls (Continued) | |
|---|---|
| Step 9 Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port.<br><br>You do not need to enable heartbeat backup if you are using the management port for the control link. | 1. In **Device > High Availability > General**, edit the Election Settings.<br>2. Select **Heartbeat Backup**.<br>To allow the heartbeats to be transmitted between the firewalls, you must verify that the management port across both peers can route to each other.<br><br>Enabling heartbeat backup also allows you to prevent a split-brain situation. Split brain occurs when the HA1 link goes down causing the firewall to miss heartbeats, although the firewall is still functioning. In such a situation, each peer believes that the other is down and attempts to start services that are running, thereby causing a split brain. When the heartbeat backup link is enabled, split brain is prevented because redundant heartbeats and hello messages are transmitted over the management port. |
| Step 10 Set the device priority and enable preemption.<br><br>This setting is only required if you want to make sure that a specific firewall is the preferred active firewall. For information, see Device Priority and Preemption. | 1. In **Device > High Availability > General**, edit the Election Settings.<br>2. Set the numerical value in **Device Priority**. Make sure to set a lower numerical value on the firewall that you want to assign a higher priority to.<br><br>If both firewalls have the same device priority value, the firewall with the lowest MAC address on the HA1 control link will become the active firewall.<br>3. Select **Preemptive**.<br>You must enable preemptive on both the active firewall and the passive firewall. |
| Step 11 (Optional) Modify the failover timers.<br><br>By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments. | 1. In **Device > High Availability > General**, edit the Election Settings.<br>2. Select the **Aggressive** profile for triggering failover faster; select **Advanced** to define custom values for triggering failover in your set up.<br><br>To view the preset value for an individual timer included in a profile, select **Advanced** and click **Load Recommended** or **Load Aggressive**. The preset values for your hardware model will be displayed on screen. |

 © Palo Alto Networks, Inc.

| Connect and Configure the Firewalls (Continued) | |
|---|---|
| Step 12 (Optional, only configured on the passive firewall) Modify the link status of the HA ports on the passive firewall.<br><br>The passive link state is **shutdown**, by default. After you enable HA, the link state for the HA ports on the active firewall will be green and those on the passive firewall will be down and display as red. | Setting the link state to **Auto** allows for reducing the amount of time it takes for the passive firewall to take over when a failover occurs and it allows you to monitor the link state.<br><br>To enable the link status on the passive firewall to stay up and reflect the cabling status on the physical interface:<br>1. In **Device > High Availability > General**, edit the Active Passive Settings.<br>2. Set the **Passive Link State** to **Auto**.<br>The auto option decreases the amount of time it takes for the passive firewall to take over when a failover occurs.<br><br>Although the interface displays green (as cabled and up) it continues to discard all traffic until a failover is triggered.<br><br>When you modify the passive link state, make sure that the adjacent devices do not forward traffic to the passive firewall based only on the link status of the firewall. |
| Step 13 Enable HA. | 1. In **Device > High Availability > General**, edit the Setup section.<br>2. Select **Enable HA**.<br>3. Select **Enable Config Sync**. This setting enables the synchronization of the configuration settings between the active and the passive firewall and should always be enabled.<br>4. Enter the IP address assigned to the control link of the peer firewall in **Peer HA1 IP Address**.<br> <br>For firewalls without dedicated HA ports, if the peer uses the management port for the HA1 link, enter the management port IP address of the peer.<br>5. Enter the **Backup HA1 IP Address**. |
| Step 14 Save your configuration changes. | Click **Commit**. |
| Step 15 Complete Step 2 through Step 14 on the other firewall in the HA pair. | |

| Connect and Configure the Firewalls (Continued) | |
|---|---|
| Step 16 After you finish configuring both firewalls, verify that the firewalls are paired in active/passive HA. | 1. Access the **Dashboard** on both firewalls, and view the High Availability widget.<br>2. On the active firewall, click the **Sync to peer** link.<br>3. Confirm that the firewalls are paired and synced, as shown below: |
| On the passive firewall: the state of the local firewall should display **passive** and the Running Config should show as **synchronized**. | On the active firewall: The state of the local firewall should display **active** and the Running Config should show as **synchronized**. |
| | |


© Palo Alto Networks, Inc.

# Define HA Failover Conditions

| Configure the Failover Triggers | | |
|---|---|---|
| Step 1 | To configure link monitoring, define the interfaces that you would like to monitor. A change in the link state of these interface will trigger a failover. | 1. Select **Device > High Availability > Link and Path Monitoring.**<br>2. In the **Link Group** section, click **Add.**<br>3. Name the **Link Group, Add** the interfaces to monitor, and select the **Failure Condition** for the group. The Link group you define is added to the **Link Group** section.<br> |
| Step 2 | (Optional) Modify the failure condition for the Link Groups that you configured (in the preceding step) on the firewall.<br>By default, the firewall will trigger a failover when any monitored link fails. | 1. Select the **Link Monitoring** section.<br>2. Set the **Failure Condition** to **All.**<br>The default setting is **Any**. |
| Step 3 | To configure path monitoring, define the destination IP addresses that the firewall should ping to verify network connectivity. | 1. In the **Path Group** section of the **Device > High Availability > Link and Path Monitoring** tab, pick the **Add option for your set up:** Virtual Wire, VLAN, or Virtual Router.<br>2. Select the appropriate item from the drop-down for the **Name** and **Add** the IP addresses (source and/or destination, as prompted) that you wish to monitor. Then select the **Failure Condition** for the group. The path group you define is added to the **Path Group** section.<br> |
| Step 4 | (Optional) Modify the failure condition for all Path Groups configured on the firewall.<br>By default, the firewall will trigger a failover when any monitored path fails. | Set the **Failure Condition** to **All.**<br>The default setting is **Any**. |
| Step 5 | Save your changes. | Click **Commit**. |

If you are using SNMPv3 to monitor the firewalls, note that the SNMPv3 Engine ID is unique to each firewall; the EngineID is not synchronized between the HA pair and, therefore, allows you to independently monitor each firewall in the HA pair. For information on setting up SNMP, see Set Up SNMP Trap Destinations.

Because the EngineID is generated using the firewall's unique serial number, on the VM-Series firewall you must apply a valid license in order to obtain a unique EngineID for each firewall.

## Verify Failover

To test that your HA configuration works properly, trigger a manual failover and verify that the firewalls transition states successfully.

| Verify Failover | | |
|---|---|---|
| Step 1 | Suspend the active firewall. | Click the **Suspend local device** link on the **Device > High Availability > Operational Commands** tab.<br> 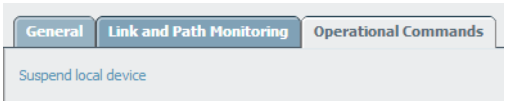 |
| Step 2 | Verify that the passive firewall has taken over as active. | On the **Dashboard**, verify that the state of the passive firewall changes to **active** in the **High Availability** widget.<br> |
| Step 3 | Restore the suspended firewall to a functional state. Wait for a couple minutes, and then verify that preemption has occurred, if preemptive is enabled. | 1. On the firewall you previously suspended, select the **Make local device functional** link on the **Device > High Availability > Operational Commands** tab.<br> 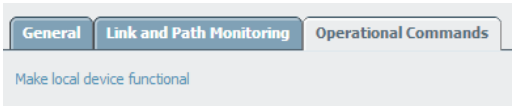 <br>2. In the **High Availability** widget on the **Dashboard,** confirm that the firewall has taken over as the active firewall and that the peer is now in a passive state.<br> |

 © Palo Alto Networks, Inc.

# HA Resources

For more information on HA, refer to the following sources:

- Active/Active HA
- High Availability Synchronization
- High Availability Failover Optimization
- Upgrading an HA pair
- Examples: Deploying HA

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

# Reports and Logging

The firewall provides reports and logs that are useful for monitoring activity on your network. You can monitor the logs and filter the information to generate reports with predefined or customized views. You can, for example, use the predefined templates to generate reports on a user's activity or analyze the reports and logs to interpret unusual behavior on your network and generate a custom report on the traffic pattern. The following topics describe how to view, manage, customize, and generate the reports and logs on the firewall:

- Use the Dashboard
- Use the Application Command Center
- Use App-Scope
- Take Packet Captures
- Monitor the Firewall
- Forward Logs to External Services
- Monitor the Firewall Using SNMP
- Monitor the Firewall Using NetFlow
- NetFlow Templates
- Identify Firewall Interfaces in External Monitoring Systems
- Manage Reporting
- Syslog Field Descriptions

# Use the Dashboard

The **Dashboard** tab widgets show general device information, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available widgets are displayed by default, but each administrator can remove and add individual widgets, as needed.

Click the refresh icon to update the Dashboard or an individual widget. To change the automatic refresh interval, select an interval from the drop-down (**1 min**, **2 mins**, **5 mins**, or **Manual**). To add a widget to the Dashboard, click the widget drop-down, select a category and then the widget name. To delete a widget, click in the title bar.

The following table describes the Dashboard widgets.

| Dashboard Charts | Descriptions |
|---|---|
| Top Applications | Displays the applications with the most sessions. The block size indicates the relative number of sessions (mouse-over the block to view the number), and the color indicates the security risk—from green (lowest) to red (highest). Click an application to view its application profile. |
| Top High Risk Applications | Similar to Top Applications, except that it displays the highest-risk applications with the most sessions. |
| General Information | Displays the device name, model, PAN-OS software version, the application, threat, and URL filtering definition versions, the current date and time, and the length of time since the last restart. |
| Interface Status | Indicates whether each interface is up (green), down (red), or in an unknown state (gray). |
| Threat Logs | Displays the threat ID, application, and date and time for the last 10 entries in the Threat log. The threat ID is a malware description or URL that violates the URL filtering profile. |
| Config Logs | Displays the administrator username, client (Web or CLI), and date and time for the last 10 entries in the Configuration log. |
| Data Filtering Logs | Displays the description and date and time for the last 60 minutes in the Data Filtering log. |
| URL Filtering Logs | Displays the description and date and time for the last 60 minutes in the URL Filtering log. |
| System Logs | Displays the description and date and time for the last 10 entries in the System log.<br> A `Config installed` entry indicates configuration changes were committed successfully. |
| System Resources | Displays the Management CPU usage, Data Plane usage, and the Session Count, which displays the number of sessions established through the firewall. |
| Logged In Admins | Displays the source IP address, session type (Web or CLI), and session start time for each administrator who is currently logged in. |

 © Palo Alto Networks, Inc.

| Dashboard Charts | Descriptions |
|---|---|
| ACC Risk Factor | Displays the average risk factor (1 to 5) for the network traffic processed over the past week. Higher values indicate higher risk. |
| High Availability | If high availability (HA) is enabled, indicates the HA status of the local and peer device—green (active), yellow (passive), or black (other). For more information about HA, see High Availability. |
| Locks | Shows configuration locks taken by administrators. |

# Use the Application Command Center

The **ACC** tab visually depicts trends and historic view of traffic on your network.


- ACC Risk Level
- ACC Charts
- ACC Detail Pages
- Use the ACC


 © Palo Alto Networks, Inc.

## ACC Risk Level

The **ACC** tab displays the overall risk level for all network traffic, the risk levels and number of threats detected for the most active and highest-risk applications on your network, and the number of threats detected from the busiest application categories and from all applications at each risk level. Use the ACC to view application data for the past hour, day, week, month, or any custom-defined time frame. **Risk** levels (1=lowest to 5=highest) indicate the application's relative security risk based on criteria such as whether the application can share files, is prone to misuse, or tries to evade firewalls.

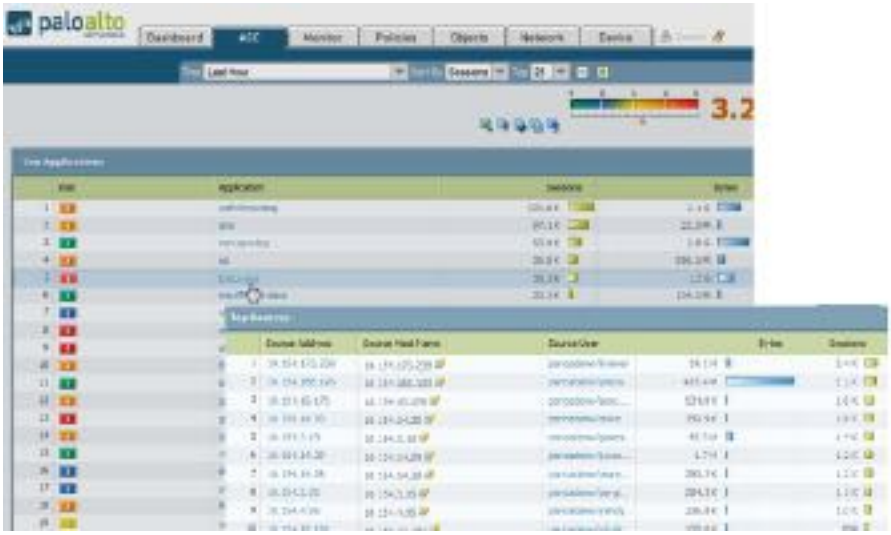

## ACC Charts

There are five charts displayed on the **Application Command Center (ACC)** tab:

- Application
- URL Filtering
- Threat Prevention
- Data Filtering
- HIP Matches

| ACC Chart | Description |
|---|---|
| Application | Displays application information grouped by the following attributes:<br>• Applications<br>• High risk applications<br>• Categories<br>• Sub Categories<br>• Technology<br>• Risk<br>Each chart can include the number of sessions, bytes transmitted and received, number of threats, application category, application subcategories, application technology, and risk level, as applicable. |
| URL Filtering | Displays URL/category information grouped by the following attributes:<br>• URL Categories<br>• URLs<br>• Blocked URL Categories<br>• Blocked URLs<br>Each chart can include the URL, URL category, repeat count (number of times access was attempted, as applicable). |
| Threat Prevention | Displays threat information grouped by the following attributes:<br>• Threats<br>• Types<br>• Spyware<br>• Spyware Phone Home<br>• Spyware Downloads<br>• Vulnerability<br>• Virus<br>Each chart can include the threat ID, count (number of occurrences), number of sessions, and subtype (such as vulnerability), as applicable. |

 © Palo Alto Networks, Inc.

| ACC Chart | Description |
|---|---|
| Data Filtering | Displays information on data filtered by the firewall grouped by the following attributes:<br>• Content/File Types<br>• Types<br>• File Names |
| HIP Matches | Displays the host information collected by the firewall grouped by:<br>• HIP Objects<br>• HIP Profiles |

## ACC Detail Pages

To view additional details, click any of the links on the ACC charts. A details page opens to show information about the item at the top and additional lists for related items. For example, click on the web-browsing link on the Application chart opens the Application Information page for web-browsing:

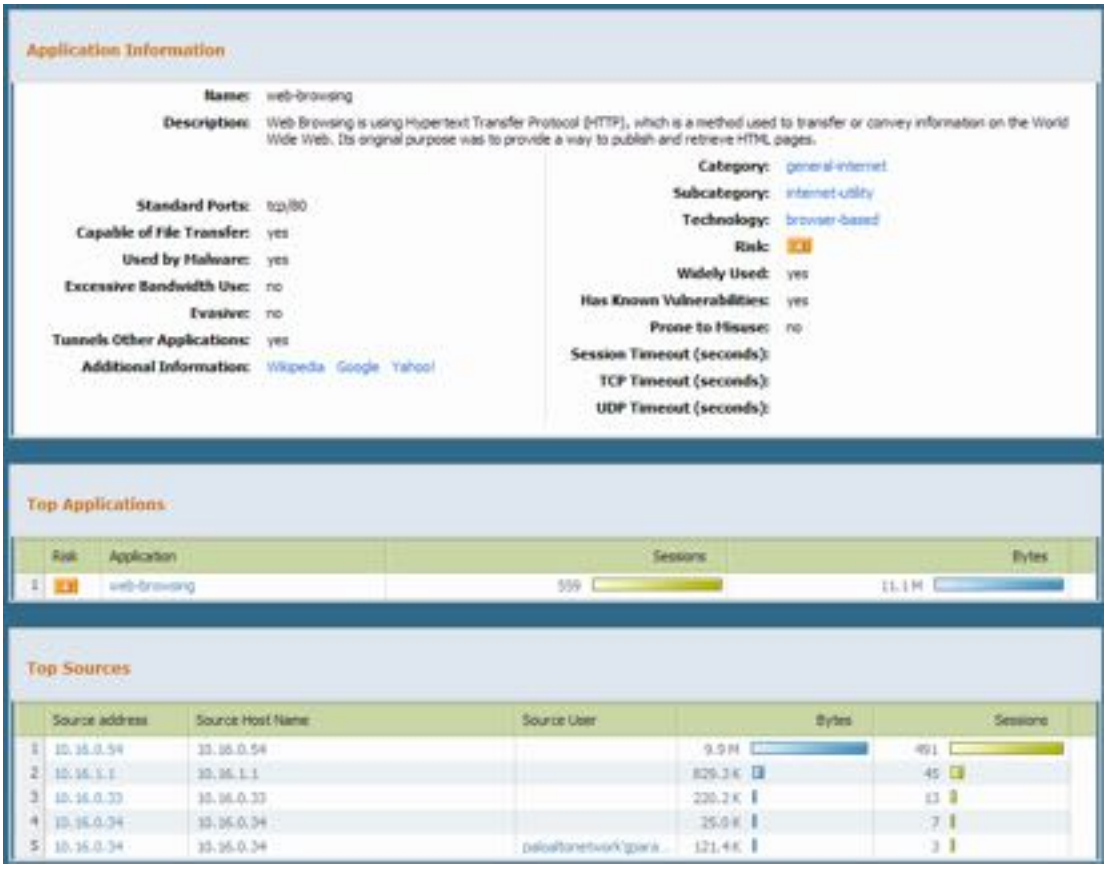


 © Palo Alto Networks, Inc.

## Use the ACC

The following procedure describes how to use the **ACC** tab and how to customize your view:

| Use the ACC | |
|---|---|
| Step 1 | On the **ACC**, change one or more of the settings at the top of the page.<br>• Use the drop-down to select Applications, URL Categories, Threats, Content/File Types, and HIP Objects to view.<br>• Select a virtual system, if virtual systems are defined.<br>• Select a time period from the **Time** drop-down. The default is Last Hour.<br>• Select a sorting method from the **Sort By** drop-down. You can sort the charts in descending order by number of sessions, bytes, or threats. The default is by number of sessions.<br>• For the selected sorting method, select the top number of applications and application categories shown in each chart from the **Top** drop-down. Click the Submit icon to apply the selected settings. |
| Step 2 | To open log pages associated with the information on the page, use the log links in the upper-right corner of the page, as shown here. The context for the logs matches the information on the page.<br> |
| Step 3 | To filter the list, click an item in one of the columns, this will add that item to the filter bar located above the log column names. After adding the desired filters, click the Apply Filter icon.<br> 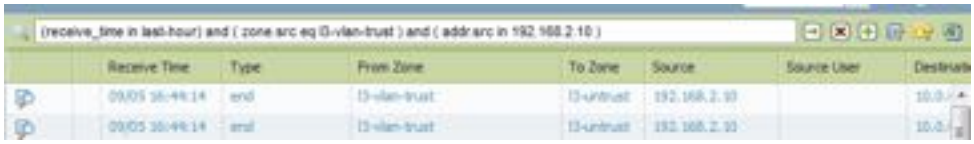 |

# Use App-Scope

The App-Scope reports provide visibility and analysis tools to help pinpoint problematic behavior, helping you understand changes in application usage and user activity, users and applications that take up most of the network bandwidth, and identify network threats.

With the App-Scope reports, you can quickly see if any behavior is unusual or unexpected. Each report provides a dynamic, user-customizable window into the network; hovering the mouse over and clicking either the lines or bars on the charts opens detailed information about the specific application, application category, user, or source on the ACC. The App-Scope charts give you the ability to:

- Toggle the attributes in the legend to only view chart details that you want to review. The ability to include or exclude data from the chart allows you to change the scale and review details more closely.
- Click into an attribute in a bar chart and drill down to the related sessions in the ACC. Click into an Application name, Application Category, Threat Name, Threat Category, Source IP address or Destination IP address on any bar chart to filter on the attribute and view the related sessions in the ACC.
- Export a chart or map to PDF or as an image. For portability and offline viewing, you can Export charts and maps as PDFs or PNG images.

The following App-Scope reports are available:

- Summary Report
- Change Monitor Report
- Threat Monitor Report
- Threat Map Report
- Network Monitor Report
- Traffic Map Report

 © Palo Alto Networks, Inc.

## Summary Report

The App-Scope Summary report displays charts for the top five gainers, losers, and bandwidth consuming applications, application categories, users, and sources.

## Change Monitor Report

The App-Scope Change Monitor report displays changes over a specified time period. For example, the following chart displays the top applications that gained in use over the last hour as compared with the last 24-hour period. The top applications are determined by session count and sorted by percent.

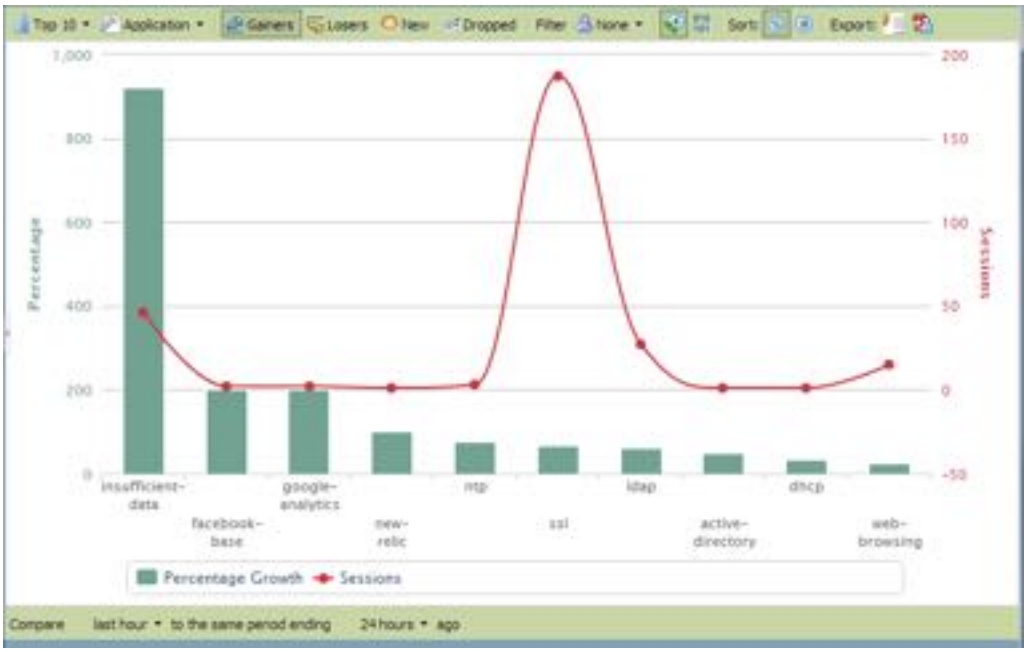


The Change Monitor Report contains the following buttons and options.

| Button | Description |
|---|---|
| Top 10 | Determines the number of records with the highest measurement included in the chart. |
| Application | Determines the type of item reported: Application, Application Category, Source, or Destination. |
| Gainers | Displays measurements of items that have increased over the measured period. |
| Losers | Displays measurements of items that have decreased over the measured period. |
| New | Displays measurements of items that were added over the measured period. |
| Dropped | Displays measurements of items that were discontinued over the measured period. |

 © Palo Alto Networks, Inc.

| Button | Description |
|---|---|
| 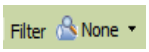 | Applies a filter to display only the selected item. **None** displays all entries. |
| | Determines whether to display session or byte information. |
| 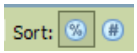 | Determines whether to sort entries by percentage or raw growth. |
| 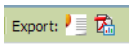 | Exports the graph as a .png image or as a PDF. |
| 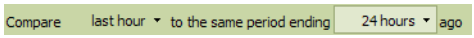 | Specifies the period over which the change measurements are taken. |

## Threat Monitor Report

The App-Scope Threat Monitor report displays a count of the top threats over the selected time period. For example, the following figure shows the top 10 threat types over the last 6 hours.

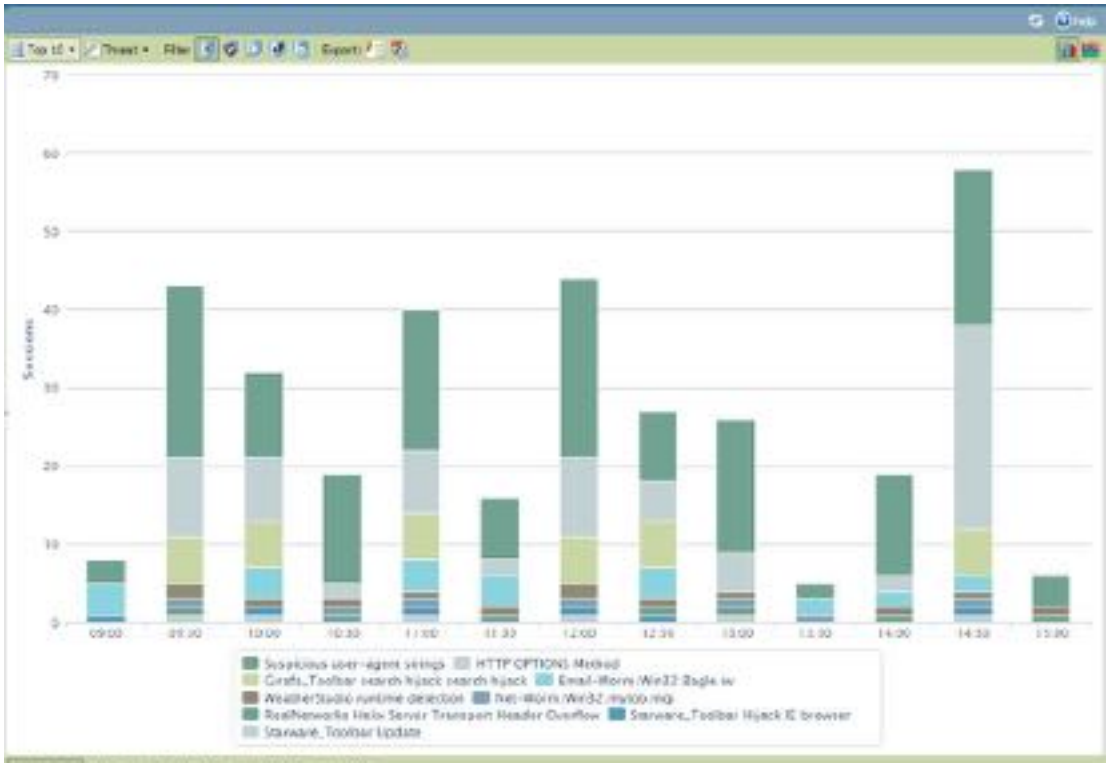

Each threat type is color-coded as indicated in the legend below the chart. The Threat Monitor report contains the following buttons and options.

| Button | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| | Determines the type of item measured: Threat, Threat Category, Source, or Destination. |
| | Applies a filter to display only the selected type of items. |
| | Determines whether the information is presented in a stacked column chart or a stacked area chart. |
| 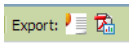 | Exports the graph as a .png image or as a PDF. |
| | Specifies the period over which the measurements are taken. |

 © Palo Alto Networks, Inc.

## Threat Map Report

The App-Scope Threat Map report shows a geographical view of threats, including severity. Each threat type is color-coded as indicated in the legend below the chart.

The firewall uses geolocation for creating threat maps. The firewall is placed at the bottom of the threat map screen, if you have not specified the geolocation coordinates (**Device > Setup > Management,** General Settings section) on the firewall.

The Threat Map report contains the following buttons and options.

| Button | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| 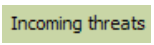 | Displays incoming threats. |
| 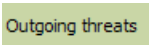 | Displays outgoing threats. |
| | Applies a filter to display only the selected type of items. |
|  | Zoom in and zoom out of the map. |
| | Exports the graph as a .png image or as a PDF. |
| | Indicates the period over which the measurements are taken. |

## Network Monitor Report

The App-Scope Network Monitor report displays the bandwidth dedicated to different network functions over the specified period of time. Each network function is color-coded as indicated in the legend below the chart. For example, the image below shows application bandwidth for the past 7 days based on session information.

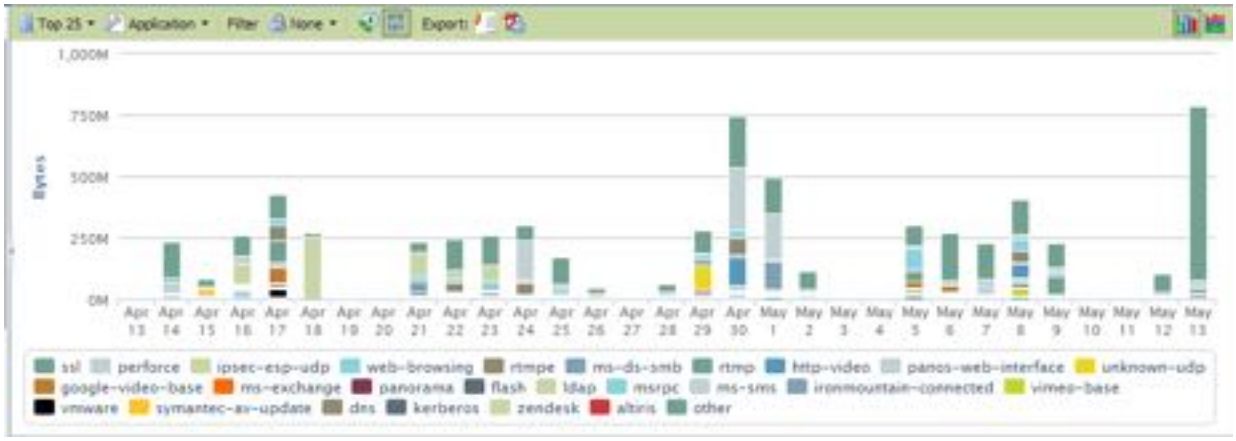


The Network Monitor report contains the following buttons and options.

| Button | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| 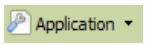  | Determines the type of item reported: Application, Application Category, Source, or Destination. |
| 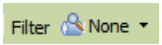  | Applies a filter to display only the selected item. **None** displays all entries. |
| | Determines whether to display session or byte information. |
| 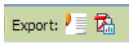  | Exports the graph as a .png image or as a PDF. |
| | Determines whether the information is presented in a stacked column chart or a stacked area chart. |
| | Indicates the period over which the change measurements are taken. |

 © Palo Alto Networks, Inc.

## Traffic Map Report

The App-Scope Traffic Map report shows a geographical view of traffic flows according to sessions or flows.

The firewall uses geolocation for creating traffic maps. The firewall is placed at the bottom of the traffic map screen, if you have not specified the geolocation coordinates (**Device > Setup > Management,** General Settings section) on the firewall.

Each traffic type is color-coded as indicated in the legend below the chart. The Traffic Map report contains the following buttons and options.

| Buttons | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| 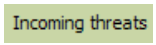 | Displays incoming threats. |
| 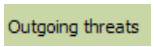 | Displays outgoing threats. |
| | Determines whether to display session or byte information. |
|  | Zoom in and zoom out of the map. |
| | Exports the graph as a .png image or as a PDF. |
| | Indicates the period over which the change measurements are taken. |

# Take Packet Captures

PAN-OS supports packet capture for troubleshooting or detecting unknown applications. You can define filters such that only the packets that match the filters are captured. The packet captures are locally stored on the device and are available for download to your local computer.

Packet Capture is for troubleshooting only. This feature can cause the system performance to degrade and should be used only when necessary. Remember to disable the feature after you complete the packet capture.

The following table describes the packet capture settings on **Monitor > Packet Capture**.

| Field | Description |
|---|---|
| **Manage Filters** | Click **Manage Filters**, click **Add** to add a new filter, and specify the following information:<br>• **Id**—Enter or select an identifier for the filter.<br>• **Ingress Interface**—Select the firewall interface.<br>• **Source**—Specify the source IP address.<br>• **Destination**—Specify the destination IP address.<br>• **Src Port**—Specify the source port.<br>• **Dest Port**—Specify the destination port.<br>• **Proto**—Specify the protocol to filter.<br>• **Non-IP**—Choose how to treat non-IP traffic (exclude all IP traffic, include all IP traffic, include only IP traffic, or do not include an IP filter).<br>• **IPv6**—Select the check box to include IPv6 packets in the filter. |
| **Filtering** | Click to toggle the filtering selections on or off. |
| **Pre-Parse Match** | Click to toggle the pre-parse match option on or off.<br>The pre-parse-match option is added for advanced troubleshooting purposes. After a packet enters the ingress port, it proceeds through several processing steps before it is parsed for matches against pre-configured filters.<br>It is possible for a packet, due to a failure, to not reach the filtering stage. This can occur, for example, if a route lookup fails.<br>Set the pre-parse-match setting to **ON** to emulate a positive match for every packet entering the system. This allows the firewall to capture even the packets that do not reach the filtering process. If a packet is able to reach the filtering stage, it is then processed according to the filter configuration and discarded if it fails to meet filtering criteria. |

 © Palo Alto Networks, Inc.

| Field | Description |
|---|---|
| **Packet Capture** | Click to toggle packet capturing on or off.<br>For anti-spyware and vulnerability protection profiles, you can enable extended packet captures for rules and exceptions defined in the profile. This functionality allows the firewall to capture from 1 to 50 packets and provides more context when analyzing the threat logs.<br>To define the extended packet capture length:<br>1. Select **Device > Setup > Content-ID.**<br>2. Edit the Threat Detection Settings section to specify the **Capture Length** for the number of packets to capture.<br> 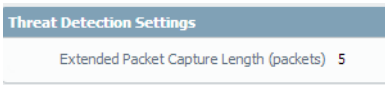 <br>3. View the packet capture in **Monitor > Logs > Threat**.<br>Locate the threat log entry and click the green arrow (Packet Capture) icon in the corresponding row to view the capture. |
| **Packet Capture Stage** | Select **Add** and specify the following:<br>• **Stage**—Indicate the point at which to capture the packet:<br>• **drop**—When packet processing encounters an error and the packet is to be dropped.<br>• **firewall**—When the packet has a session match or a first packet with a session is successfully created.<br>• **receive**—When the packet is received on the dataplane processor.<br>• **transmit**—When the packet is to be transmitted on the dataplane processor.<br>• **File**—Specify the capture file name. The file name should begin with a letter and can include letters, digits, periods, underscores, or hyphens.<br>• **Packet Count**—Specify the number of packets after which capturing stops.<br>• **Byte Count**—Specify the number of bytes after which capturing stops. |
| **Captured Files** | Select **Delete** to remove a packet capture file from the list displaying captured files. |
| **Clear All Settings** | Select **Clear All Settings** to clear all packet capture settings. |

# Monitor the Firewall

The following sections describe the methods you can use to monitor the firewall and provide basic setup instructions:

- Monitor Applications and Threats
- Monitor Log Data
- Monitor the Dashboard
- View Reports

You can also configure the firewall (excluding PA-4000 Series and PA-7050 firewalls) to export flow data to a NetFlow collector for analysis and reporting. To configure NetFlow Settings, refer to the PAN-OS Web Interface Reference Guide.


© Palo Alto Networks, Inc.

## Monitor Applications and Threats

All Palo Alto Networks next-generation firewalls come equipped with the App-ID technology, which identifies the applications traversing your network, irrespective of protocol, encryption, or evasive tactic. You can then Use the Application Command Center to monitor the applications. ACC graphically summarizes the log database to highlight the applications traversing your network, who is using them, and their potential security impact. ACC is dynamically updated, using the continuous traffic classification that App-ID performs; if an application changes ports or behavior, App-ID continues to see the traffic, displaying the results in ACC.

You can quickly investigate new, risky, or unfamiliar applications that appear in ACC with a single click that displays a description of the application, its key features, its behavioral characteristics, and who is using it. Additional visibility into URL categories, threats, and data provides a complete and well-rounded picture of network activity. With ACC, you can very quickly learn more about the traffic traversing the network and then translate that information into a more informed security policy.

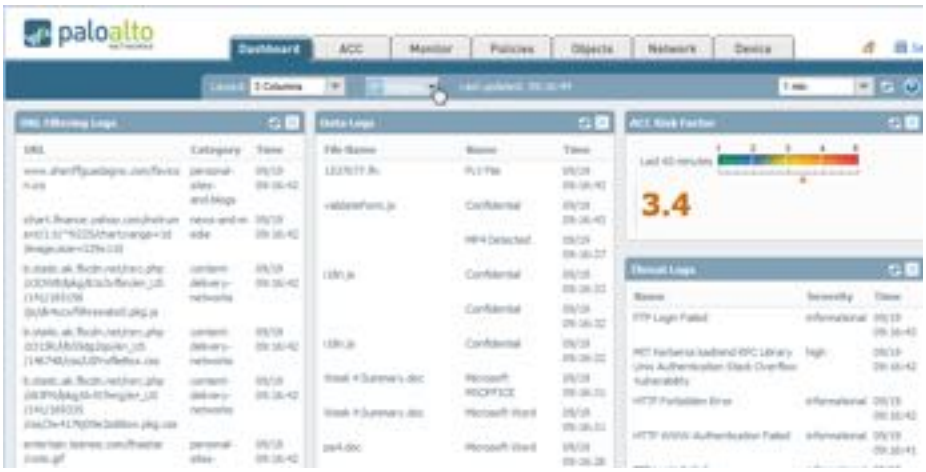

# Monitor Log Data

All Palo Alto Networks next-generation firewalls can generate log files that provide an audit trail of the activities and events on the firewall. There are separate logs for separate types of activities and events. For example, the Threat logs record all traffic that causes the firewall to generate a security alarm, URL Filtering logs record all traffic that matches a URL Filtering profile attached to a security policy, and Config logs record all changes to the firewall configuration.

You can either Forward Logs to External Services or you can view logs locally on the device as follows:

- View the Log Files
- Filter Log Data

## View the Log Files

The firewall maintains logs for WildFire, configurations, system, alarms, traffic flows, threats, URL filtering, data filtering, and Host Information Profile (HIP) matches. You can view the current logs at any time. To locate specific entries, you can apply filters to most of the log fields.

The firewall displays the information in logs so that role-based administration permissions are respected. When you display logs, only the information that you have permission to see is included. For information on administrator permissions, see Administrative Roles.

By default all log files are generated and stored locally on the firewall. You can view these log files directly (**Monitor > Logs**):

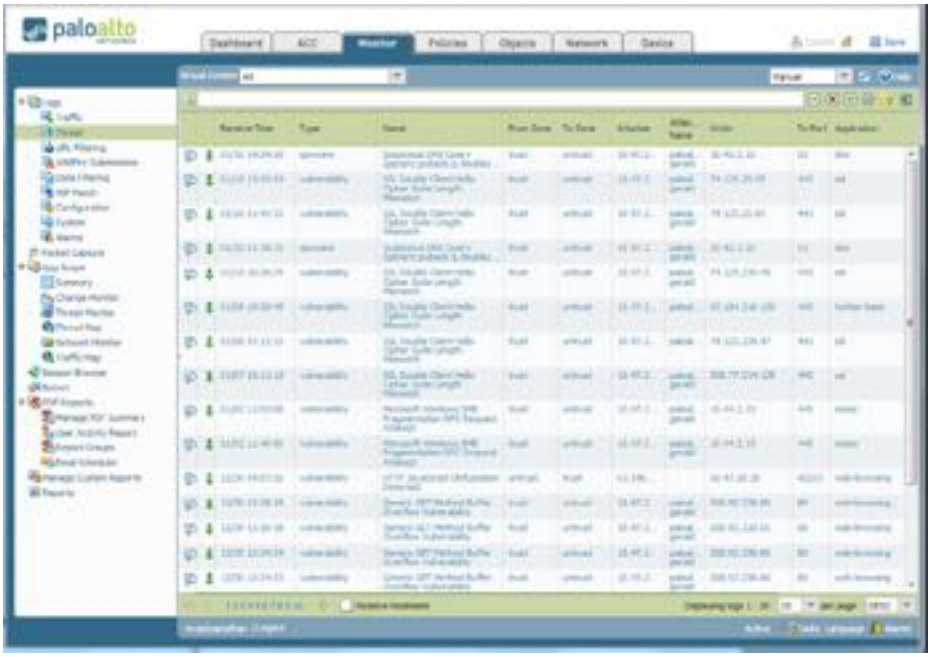

To display additional details, click the spyglass icon for an entry.

 © Palo Alto Networks, Inc.

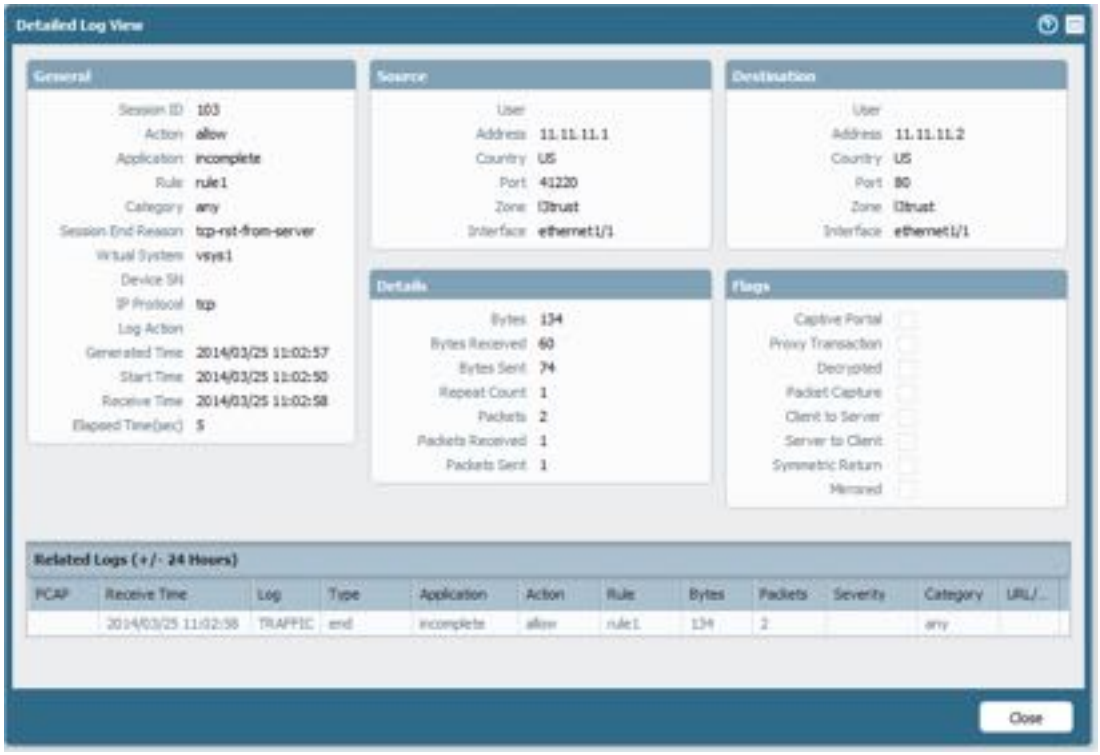


The following table includes information on each log type:

| Log Description Charts | Description |
|---|---|
| Traffic | Displays an entry for the start and end of each session. Each entry includes the date and time, source and destination zones, addresses and ports, application name, security rule name applied to the flow, rule action (**allow**, **deny**, or **drop**), ingress and egress interface, number of bytes, and session end reason.<br>Click next to an entry to view additional details about the session, such as whether an ICMP entry aggregates multiple sessions between the same source and destination (the **Count** value will be greater than one).<br>The **Type** column indicates whether the entry is for the start or end of the session, or whether the session was denied or dropped. A **drop** indicates that the security rule that blocked the traffic specified **any** application, while a **deny** indicates the rule identified a specific application.<br>If traffic is dropped before the application is identified, such as when a rule drops all traffic for a specific service, the application is shown as **not-applicable**. |

| Log Description Charts | Description |
|---|---|
| Threat | Displays an entry when traffic matches a Security Profile (Antivirus, Anti-spyware, Vulnerability, URL Filtering, File Blocking, Data Filtering, or DoS Protection) that is attached to a security policy on the firewall. Each entry includes the date and time, a threat name or URL, the source and destination zones, addresses, and ports, the application name, and the alarm action (**allow** or **block**) and severity.<br>Click next to an entry to view additional details about the threat, such as whether the entry aggregates multiple threats of the same type between the same source and destination (the **Count** value will be greater than one).<br>The **Type** column indicates the type of threat, such as "virus" or "spyware." The Name column is the threat description or URL, and the **Category** column is the threat category (such as "keylogger") or URL category.<br>If local packet captures are enabled, click next to an entry to access the captured packets. To enable local packet captures, see Take Packet Captures. |
| URL Filtering | Displays logs for all traffic that matches a URL Filtering profile attached to a security policy. For example, if policy blocks access to specific web sites and web site categories or if policy is configured to generate an alert when a web site is accessed. For information on defining URL filtering profiles, see URL Filtering. |
| WildFire Submissions | Displays logs for files that are uploaded and analyzed by the WildFire cloud; log data is sent back to the device after analysis, along with the analysis results. |
| Data Filtering | Displays logs for the security policies that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall. See Set Up Data Filtering for information on defining data filtering profiles.<br>This log also shows information for file-blocking profiles. For example, if you are blocking *.exe* files, the log will show the files that were blocked. If you forward files to WildFire, you will see the results of that action. In this case, if you are forwarding PE files to WildFire, for example, the log will show that the file was forwarded and will also show the status on whether or not it was uploaded to WildFire successfully. |
| Configuration | Displays an entry for each configuration change. Each entry includes the date and time, the administrator username, the IP address from where the change was made, the type of client (XML, Web or CLI), the type of command executed, whether the command succeeded or failed, the configuration path, and the values before and after the change. |
| System | Displays an entry for each system event. Each entry includes the date and time, the event severity, and an event description. |
| HIP Match | Displays traffic flows that match a HIP Object or HIP Profile that you have configured. |

## Filter Log Data

Each log page has a filter area at the top of the page.

Use the filter area as follows:

 © Palo Alto Networks, Inc.

- Click any of the underlined links in the log listing to add that item as a log filter option. For example, if you click the **Host** link in the log entry for 10.0.0.252 and **Web Browsing**, both items are added, and the search will find entries that match both (AND search).
- To define other search criteria, click **Add Log Filter**. Select the type of search (and/or), the attribute to include in the search, the matching operator, and the values for the match, if appropriate. Click **Add** to add the criterion to the filter area on the log page, and then click **Close** to close the pop-up window. Click **Apply Filter** to display the filtered list.

If the **Value** string matches an **Operator** (such as **has** or **in**), enclose the string in quotation marks to avoid a syntax error. For example, if you filter by destination country and use `IN` as a **Value** to specify INDIA, enter the filter as `( dstloc eq "IN" )`.

You can combine filter expressions added on the log page with those you define in the Add Log Filter dialog. The filter field on the log page displays each filter as an entry.

If you add a **Receive Time** filter with the **Operator** set to **in** and the **Value** set to **Last 60 seconds**, some of the page links on the log viewer might not show results because the number of pages might grow or shrink due to the dynamic nature of the selected time.

- To clear filters and redisplay the unfiltered list, click **Clear Filter**.
- To save your selections as a new filter, click **Save Filter**, enter a name for the filter, and click **OK**.
- To export the current log listing (as shown on the page, including any applied filters) click **Save Filter**. Select whether to open the file or save it to disk, and select the check box if you want to always use the same option. Click **OK**.
- To export the current log listing in CSV Format, select the **Export to CSV** icon. By default, exporting the log listing to CSV format generates a CSV report with up to 2,000 rows of logs. To change the limit for rows displayed in CSV reports, use the **Max Rows in CSV Export** field on the **Log Export and Reporting** tab (select **Device > Setup > Management > Logging and Reporting Settings**).

To change the automatic refresh interval, select an interval from the drop-down (**1 min**, **30 seconds**, **10 seconds**, or **Manual**).

To change the number of log entries per page, select the number of rows from the **Rows** drop-down.

Log entries are retrieved in blocks of 10 pages. Use the paging controls at the bottom of the page to navigate through the log list. Select the **Resolve Hostname** check box to begin resolving external IP addresses to domain names.

## Monitor the Dashboard

You can also monitor the local log data directly from the **Dashboard** by adding the associated widgets:

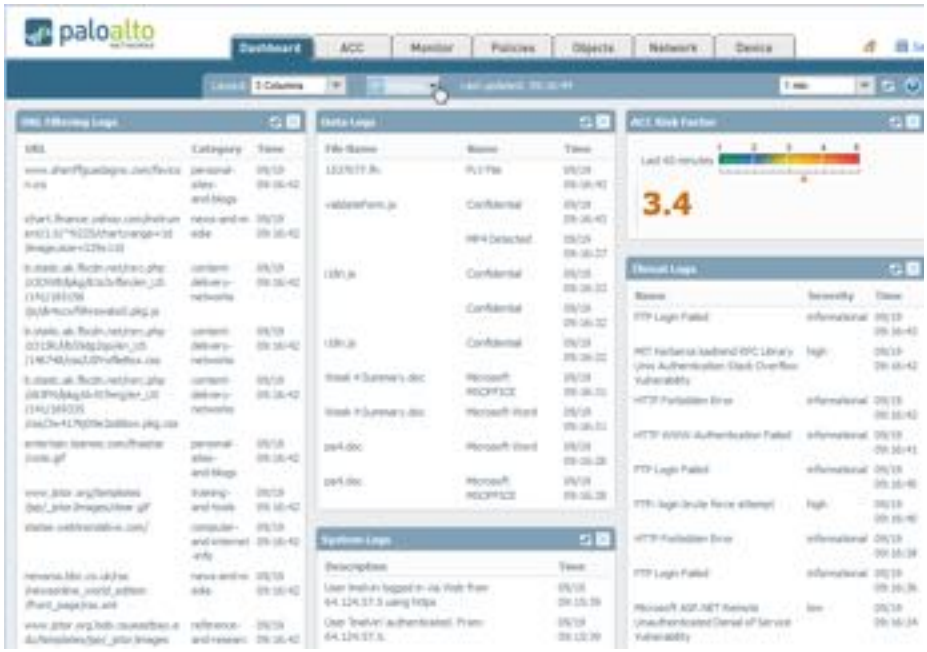


 © Palo Alto Networks, Inc.

## View Reports

The firewall also uses the log data to generate reports (**Monitor > Reports**) that display the log data in a tabular or graphical format. See About Reports for more details on the predefined and custom reports available on the firewall.

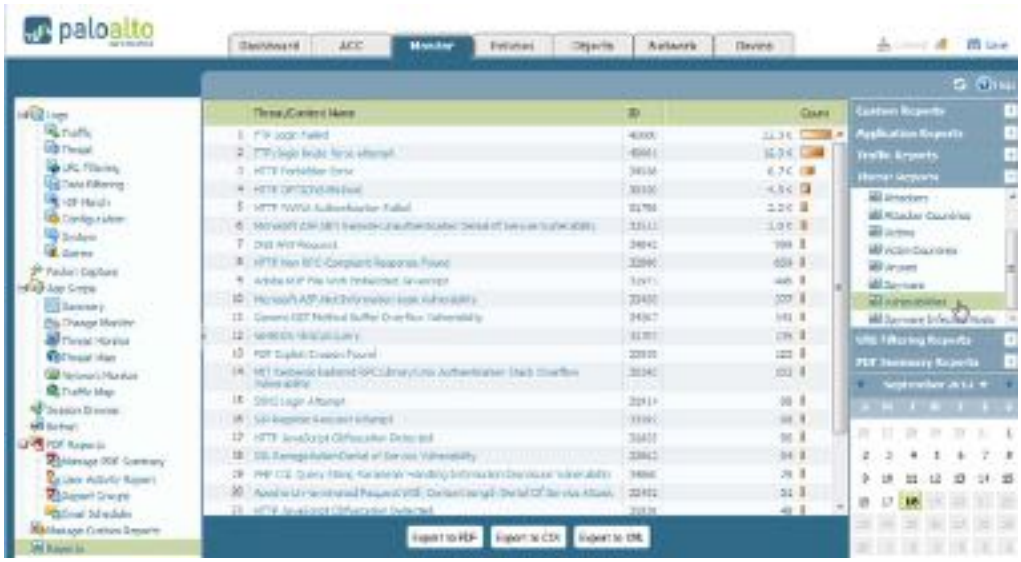

# Forward Logs to External Services

Depending on the type and severity of the data in the log files, you may want to be alerted to critical events that require your attention, or you may have policies that require you to archive the data for longer than it can be stored on the firewall. In these cases you will want to forward your log data to an external service for archive, notification, and/or analysis.

To forward log data to an external service you must complete the following tasks:

- Configure the firewall to access the remote services that will be receiving the logs. See Define Remote Logging Destinations.
- Configure each log type for forwarding. See Enable Log Forwarding.
- For traffic and threat logs, enabling log forwarding includes setting up a log forwarding profile or a default log forwarding profile. For details, see Log Forwarding Profiles.

 © Palo Alto Networks, Inc.

## Define Remote Logging Destinations

In order to reach an external service—such as a Syslog server or SNMP trap manager—the firewall must know the details of how to access and, if necessary, authenticate to the service. On the firewall, you define this information in a Server Profile. You must create a Server Profile for each external service you want the firewall to interact with. The type of logging destinations you need to set up and which logs you forward will depend on your needs. Some common log forwarding scenarios include the following:

- For immediate notification about critical system events or threats that require your attention, you can generate SNMP traps or send email alerts. See Set Up Email Alerts and/or Set Up SNMP Trap Destinations.
- For long-term storage and archival of data and for centralized device monitoring, you can send the log data to a Syslog server. See Define Syslog Servers. This enables integration with third-party security monitoring tools, such as Splunk! or ArcSight. You can also secure the channel between the firewall and the Syslog server. See Configure the Firewall to Authenticate to the Syslog Server.
- For aggregation and reporting of log data from multiple Palo Alto Networks firewalls, you can forward logs to a Panorama Manager or Panorama Log Collector. See Enable Log Forwarding.

You can define as many Server Profiles as you need. For example, you could use separate Server Profiles to send traffic logs to one Syslog server and system logs to a different one. Or, you could include multiple server entries in a single Server Profile to enable you to log to multiple Syslog servers for redundancy.

By default, all log data is forwarded over the MGT interface. If you plan to use an interface other than MGT, you will need to configure a Service Route for each service to which you plan to forward logs as described in Step 5 in the procedure to Set Up Network Access for External Services.

## Set Up Email Alerts

| Set Up Email Alerts | | |
|---|---|---|
| Step 1 | Create a Server Profile for your email server. | 1. Select **Device > Server Profiles > Email**.<br>2. Click **Add** and then enter a **Name** for the profile.<br>3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.<br>4. Click **Add** to add a new email server entry and enter the information required to connect to the Simple Mail Transport Protocol (SMTP) server and send email (you can add up to four email servers to the profile):<br>• **Server**—Name to identify the mail server (1-31 characters). This field is just a label and does not have to be the host name of an existing SMTP server.<br>• **Display Name**—The name to show in the From field of the email.<br>• **From**—The email address where notification emails will be sent from.<br>• **To**—The email address to which notification emails will be sent.<br>• **Additional Recipient**—If you want the notifications sent to a second account, enter the additional address here. You can only add one additional recipient. To add multiple recipients, add the email address of a distribution list.<br>• **Gateway**—The IP address or host name of the SMTP gateway to use to send the emails.<br>5. Click **OK** to save the server profile. |
| Step 2 | (Optional) Customize the format of the email messages the firewall sends. | Select the **Custom Log Format** tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide. |
| Step 3 | Save the server profile and commit your changes. | 1. Click **OK** to save the profile.<br>2. Click **Commit** to save the changes to the running configuration. |

## Set Up SNMP Trap Destinations

Simple Network Management Protocol (SNMP) is a standard facility for monitoring the devices on your network. You can configure the firewall to send SNMP traps to your SNMP management software to alert you to critical system events or threats that require your immediate attention.

You can also use SNMP to monitor the firewall. In this case, your SNMP manager must be configured to get statistics from the firewall rather than (or in addition to) having the firewall send traps to the manager. For more information, see Configure the Firewall to Authenticate to the Syslog Server.

 © Palo Alto Networks, Inc.

**Set Up SNMP Trap Destinations**

Step 1 (SNMP v3 only) Get the engine ID for the firewall.

In many cases, the MIB browser or SNMP manager will automatically discover the engine ID upon successful connection to the SNMP agent on the firewall. You can usually find this information in the agent settings section of the interface. Refer to the documentation for your specific product for instructions on finding the agent information.

In order to find out the firewall's engine ID, you must configure the firewall for SNMP v3 and send a GET message from your SNMP manager or MIB browser as follows:

1. Enable the interface to allow inbound SNMP requests:
   - If you will be receiving SNMP GET messages on the MGT interface, select **Device > Setup > Management** and click to Edit in the Management Interface Settings section of the screen. In the **Services** section, select the **SNMP** check box and then click **OK**.
   - If you will be receiving SNMP GET messages on a different interface, you must associate a management profile with the interface and enable SNMP management.
2. Configure the firewall for SNMP v3 as described in Step 2 in Set Up SNMP Monitoring. If you do not configure the firewall for SNMP v3, your MIB browser will not allow you to GET the engine ID.
3. Connect your MIB browser or SNMP manager to the firewall and run a GET for OID 1.3.6.1.6.3.10.2.1.1.0. The value that is returned is the unique engine ID for the firewall.

Step 2 Create a Server Profile that contains the information for connecting and authenticating to the SNMP manager(s).

1. Select **Device > Server Profiles > SNMP Trap**.
2. Click **Add** and then enter a **Name** for the profile.
3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.
4. Specify the version of SNMP you are using (V2c or V3).
5. Click **Add** to add a new **SNMP Trap Receiver** entry (you can add up to four trap receivers per server profile). The required values depend on whether you are using SNMP V2c or V3 as follows:

   SNMP V2c

   - **Server**—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the host name of an existing SNMP server.
   - **Manager**—The IP address of the SNMP manager to which you want to send traps.
   - **Community**—The community string required to authenticate to the SNMP manager.

   SNMP V3

   - **Server**—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the host name of an existing SNMP server.
   - **Manager**—The IP address of the SNMP manager to which you want to sent traps.
   - **User**—The username required to authenticate to the SNMP manager.
   - **EngineID**—The engine ID of the firewall, as identified in Step 1. This is a hexadecimal value from 5 to 64 bytes with a 0x prefix. Each firewall has a unique engine ID.
   - **Auth Password**—The password to be used for authNoPriv level messages to the SNMP manager. This password will be hashed using Secure Hash Algorithm (SHA-1), but will not be encrypted.
   - **Priv Password**—The password to be used for authPriv level messages to the SNMP manager. This password be hashed using SHA and will be encrypted using Advanced Encryption Standard (AES 128).
6. Click **OK** to save the server profile.

| Set Up SNMP Trap Destinations (Continued) | | |
|---|---|---|
| Step 3 | (Optional) Set up a service route for SNMP traps. | By default, SNMP traps are sent over the MGT interface. If you want to use a different interface for SNMP traps, you must edit the service route to enable the firewall to reach your SNMP manager. See Set Up Network Access for External Services for instructions. |
| Step 4 | Commit your changes. | Click **Commit**. The device may take up to 90 seconds to save your changes. |
| Step 5 | Enable the SNMP manager to interpret the traps it receives from the firewall. | Load the PAN-OS MIB files into your SNMP management software and compile them. Refer to the documentation for your SNMP manager for specific instructions on how to do this. |

## Define Syslog Servers

Syslog is a standard log transport mechanism that enables the aggregation of log data from different network devices—such as routers, firewalls, printers—from different vendors into a central repository for archive, analysis, and reporting.

The firewall generates six types of logs that can be forwarded to an external syslog server: traffic, threat, WildFire, host information profile (HIP) match, config, and system. If you want to forward all or some of these logs to an external service for long-term storage and analysis, you can use TCP or SSL for reliable and secure transport of logs, or UDP for non-secure transport.

 © Palo Alto Networks, Inc.

| Set Up Syslog Forwarding | |
|---|---|
| Step 1 Create a Server Profile that contains the information for connecting to the Syslog server(s). | 1. Select **Device > Server Profiles > Syslog**.<br>2. Click **Add** and then enter a **Name** for the profile.<br>3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.<br>4. Click **Add** to add a new Syslog server entry and enter the information required to connect to the Syslog server (you can add up to four Syslog servers to the same profile):<br>• **Name**—Unique name for the server profile.<br>• **Server**—IP address or fully qualified domain name (FQDN) of the Syslog server.<br>• **Transport**—Select TCP, UDP, or SSL as the method of communication with the syslog server.<br>• **Port**—The port number on which to send Syslog messages (default is UDP on port 514); you must use the same port number on the firewall and the Syslog server.<br>• **Format**—Select the Syslog message format to use, BSD or IETF. Traditionally, BSD format is over UDP and IETF format is over TCP/SSL. For setting up secure syslog forwarding with client authentication, see Configure the Firewall to Authenticate to the Syslog Server.<br>• **Facility**—Select one of the Syslog standard values, which is used to calculate the priority (PRI) field in your Syslog server implementation. You should select the value that maps to how you use the PRI field to manage your Syslog messages.<br>5. (Optional) To customize the format of the Syslog messages the firewall sends, select the **Custom Log Format** tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide.<br>6. Click **OK** to save the server profile. |

| Set Up Syslog Forwarding (Continued) | |
|---|---|
| Step 2 (Optional) Configure the header format used in Syslog messages. Choosing the header format offers more flexibility in filtering and reporting on the log data for some SIEMs.<br> 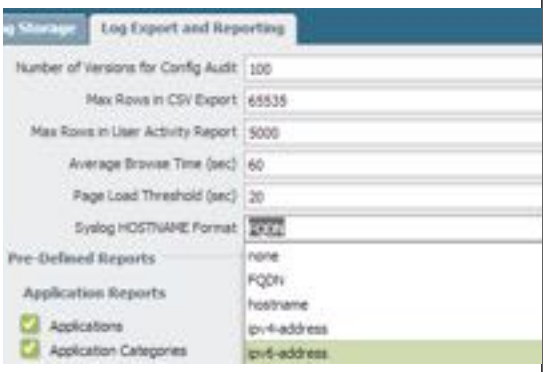 <br><br>This is a global setting and applies to all syslog server profiles configured on the appliance. | 1. Select **Device > Setup > Management** and click the Edit icon in the Logging and Reporting Settings section.<br>2. Select **Log Export and Reporting**.<br>3. Select one of the following options from the **Send Hostname in Syslog** drop-down:<br>• **FQDN**—(the default) Concatenates the hostname and domain name defined on the sending device.<br>• **hostname**—Uses the hostname defined on the sending device.<br>• **ipv4-address**—Uses the IPv4 address of the interface used to send logs on the device. By default, this is the MGT interface of the device.<br>• **ipv6-address**—Uses the IPv6 address of the interface used to send logs on the device. By default, this is the MGT interface of the device.<br>• **none**—Leaves the hostname field unconfigured on the device. There is no identifier for the device that sent the logs.<br>4. Click **OK** and **Commit**. |
| Step 3 Commit your changes. | Click **Commit**. The device may take up to 90 seconds to save your changes. |
| Step 4 Enable log forwarding. | See Enable Log Forwarding.<br>You must configure each log type for forwarding and specify the severity for which the event is logged.<br>WildFire logs are a type of threat log, but they are not logged and forwarded along with threat logs. While WildFire logs use the same syslog format as threat logs, the threat subtype is preset to WildFire. Therefore, you must enable logging/forwarding for WildFire logs distinctly from threat logs. |
| Step 5 Review the logs on the syslog server. | To parse the logs, see Syslog Field Descriptions. |

## Configure the Firewall to Authenticate to the Syslog Server

To enable client authentication for syslog over SSL, you can use a trusted CA or a self-signed CA for generating certificates that can be used for secure syslog communication. Check for the following when generating a certificate for secure syslog communication:

- The private key must be available on the sending device; the keys cannot be stored on a Hardware Security Module (HSM).
- The subject and the issuer for the certificate must not be identical.

 © Palo Alto Networks, Inc.

- The certificate is neither a trusted CA nor a certificate signing request (CSR). Neither of these types of certificates can be enabled for secure syslog communication.

**Configure the Firewall to Authenticate to the Syslog Server**

Step 1 If the syslog server requires client authentication, generate the certificate for secure communication. For details on certificates, see Certificate Management.

To verify that the sending device is authorized to communicate with the syslog server, you must enable the following:

- The server and the sending device must have certificates that are signed by the Enterprise CA; or you can generate a self-signed certificate on the firewall, export the CA root certificate from the firewall and import it in to the syslog server.
- Use the Enterprise CA or the self-signed certificate to generate a certificate with the IP address of the sending device (as the Common Name) and enabled for use in secure syslog communication. The syslog server uses this certificate to verify that the firewall is authorized to communicate with the syslog server.

  Use the following steps to generate the certificate on the firewall or Panorama:

1. Select **Device > Certificate Management > Certificates > Device Certificates**.
2. Click **Generate** to create a new certificate that will be signed by a trusted CA or the self-signed CA.
3. Enter a **Name** for the certificate.
4. In **Common Name**, enter the IP address of the device sending logs to the syslog server.
5. Select **Shared** if you want the certificate to be a shared certificate on Panorama or to be shared by all virtual systems in a multiple virtual system firewall.
6. In **Signed by**, select the trusted CA or the self-signed CA that is trusted by both the syslog server and the sending device.
7. Click **Generate**. The certificate and the keypair will be generated.
8. Click the link with name of the certificate and enable the **Certificate for Secure Syslog** option for secure access to the syslog server.
9. **Commit** the changes.
10. Verify the certificate details and that it is marked for **Usage** as **Certificate for Secure Syslog**.

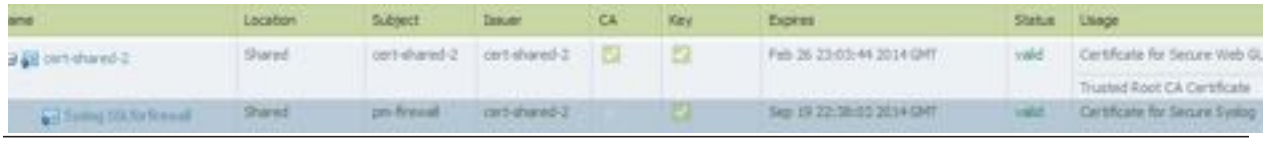

## Enable Log Forwarding

After you create the Server Profiles that define where to send your logs (see Define Remote Logging Destinations), you must enable log forwarding. For each log type, you can specify whether to forward it to Syslog, email, SNMP trap receiver, and/or Panorama.

Before you can forward log files to a Panorama Manager or a Panorama Log Collector, the firewall must be configured as a managed device. You can then enable log forwarding to Panorama for each type of log. For logs forwarded to Panorama, support for centralized log forwarding to an external syslog server is available.

You can use Secure Copy (SCP) commands from the CLI to export the entire log database (logdb) to an SCP server and import it to another firewall: refer to the PAN-OS Command Line Interface (CLI) Reference Guide. Because the log database is too large for an export or import to be practical on the following platforms, they do not support these options: PA-7050 firewall (all PAN-OS releases), Panorama virtual appliance running Panorama 6.0 or later releases, and Panorama M-Series appliances (all Panorama releases).

The way you enable forwarding depends on the log type:

- **Traffic Logs**—You enable forwarding of Traffic logs by creating a Log Forwarding Profile (**Objects > Log Forwarding**) and adding it to the security policies you want to trigger the log forwarding. Only traffic that matches a specific rule within the security policy will be logged and forwarded. For details on setting up a log forwarding profile, see Log Forwarding Profiles.
- **Threat Logs**—You enable forwarding of Threat logs by creating a Log Forwarding Profile (**Objects > Log Forwarding**) that specifies which severity levels you want to forward and then adding it to the security policies for which you want to trigger the log forwarding. A Threat log entry will only be created (and therefore forwarded) if the associated traffic matches a Security Profile (Antivirus, Anti-spyware, Vulnerability, URL Filtering, File Blocking, Data Filtering, or DoS Protection). For details on setting up a log forwarding profile, see Log Forwarding Profiles. The following table summarizes the threat severity levels:

| Severity | Description |
|---|---|
| **Critical** | Serious threats, such as those that affect default installations of widely deployed software, result in root compromise of servers, and the exploit code is widely available to attackers. The attacker usually does not need any special authentication credentials or knowledge about the individual victims and the target does not need to be manipulated into performing any special functions. |
| **High** | Threats that have the ability to become critical but have mitigating factors; for example, they may be difficult to exploit, do not result in elevated privileges, or do not have a large victim pool. |
| **Medium** | Minor threats in which impact is minimized, such as DoS attacks that do not compromise the target or exploits that require an attacker to reside on the same LAN as the victim, affect only non-standard configurations or obscure applications, or provide very limited access. In addition, WildFire Submissions log entries with a malware verdict are logged as Medium. |
| **Low** | Warning-level threats that have very little impact on an organization's infrastructure. They usually require local or physical system access and may often result in victim privacy or DoS issues and information leakage. Data Filtering profile matches are logged as Low. |

 © Palo Alto Networks, Inc.

| Severity | Description |
|---|---|
| **Informational** | Suspicious events that do not pose an immediate threat, but that are reported to call attention to deeper problems that could possibly exist. URL Filtering log entries and WildFire Submissions log entries with a benign verdict are logged as Informational. |

- **Config Logs**—You enable forwarding of Config logs by specifying a Server Profile in the log settings configuration. (**Device > Log Settings > Config Logs**).
- **System Logs**—You enable forwarding of System logs by specifying a Server Profile in the log settings configuration. (**Device > Log Settings > System Logs**). You must select a Server Profile for each severity level you want to forward. For a partial list of system log messages and their corresponding severity levels, refer to the System Log Reference. The following table summarizes the system log severity levels:

| Severity | Description |
|---|---|
| **Critical** | Hardware failures, including HA failover and link failures. |
| **High** | Serious issues, including dropped connections with external devices, such as LDAP and RADIUS servers. |
| **Medium** | Mid-level notifications, such as antivirus package upgrades. |
| **Low** | Minor severity notifications, such as user password changes. |
| **Informational** | Log in/log off, administrator name or password change, any configuration change, and all other events not covered by the other severity levels. |

## Log Forwarding Profiles

Log forwarding profiles allow you to forward traffic and threat logs to Panorama or an external system. A log forwarding profile can be added to a security zone to forward zone protection logs or to a security policy to forward logs for traffic that matches that policy. You can also configure a default log forwarding profile—the settings in the default profile will be used as the default log forwarding settings for new security zones and new security policies. This allows you to consistently include your organization's preferred log forwarding settings in new policies and zones automatically, without administrators having to manually add them each time.

The following sections show how to create a log forwarding profile and how to enable a profile to be used as the default log forwarding settings for new security policies or security zones:

- ▲ Create a Log Forwarding Profile
- ▲ Set Up or Override a Default Log Forwarding Profile

## Create a Log Forwarding Profile

Create a log forwarding profile that can be added to security policies and security zones, in order to forward traffic and threat logs to Panorama or an external system. Forwarded logs can be sent as SNMP traps, syslog messages, or email notifications.

| **Enable a Log Forwarding Profile** | | |
|---|---|---|
| Step 1 | Add a log forwarding profile. | 1. Select **Objects > Log Forwarding Profile** and **Add** a new security profile group.<br>2. Give the profile group a descriptive **Name** to help identify it when adding the profile to security policies or security zones.<br>3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.<br>4. Add settings for the Traffic logs, Threat logs, and WildFire logs:<br>• Select the **Panorama** check box for the severity of the Traffic, Threat, or WildFire logs that you want to be forwarded to Panorama.<br>• Specify logs that you want to forward to additional destinations: **SNMP Trap** destinations, **Email** servers, or **Syslog** servers.<br>5. Click **OK** to save the log forwarding profile. |
| Step 1 | Add the log forwarding profile to a security policy.<br>Traffic that matches the security policy and your log forwarding profile will be forwarded to the destinations defined in your profile.<br>For more details on security policies, see Security Policy. | 1. Select **Policies > Security** and **Add** or modify a security policy.<br>2. Select **Actions** and the log forwarding profile you created in the **Log Forwarding Profile** drop-down.<br>Threat log entries are generated according to the security profiles you have configured, in addition to the settings defined in the log forwarding profile. For more details on security profiles, see Security Profiles.<br>3. Click **OK** to save the security policy. |
| Step 1 | Add the log forwarding profile to a security zone.<br>For more details on setting up security zones, see Configure Interfaces and Zones. | 1. Select **Network > Zones** and **Add** or modify a security zone.<br>2. Select the log forwarding profile from the **Log Setting** drop-down.<br>3. Click **OK** to save the security zone. |
| Step 1 | Save your changes. | **Commit**. |

## Set Up or Override a Default Log Forwarding Profile

Add a new log forwarding profile or modify an existing one to be used as the default log forwarding settings for new security policy rules or new security zones. When an administrator creates a new security policy or a new security zone, the *default* log forwarding profile will be automatically selected as the policy or zone's log forwarding settings (the administrator can choose to manually select different log forwarding settings if desired). Use the following options to set up a default log forwarding profile or to override your default settings.

 © Palo Alto Networks, Inc.

If no default security profile exists, the profile settings for a new security policy are set to **None** by default.

| Set Up or Override a Default Log Forwarding Profile | |
|---|---|
| • Set up a default log forwarding profile. | 1. Select **Objects > Log Forwarding** and **Add** a new log forwarding profile or modify an existing profile.<br>2. **Name** the security profile group *default*:<br> 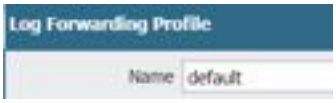 <br><br>3. Click **OK** and **Commit**.<br>4. Confirm that the *default* log forwarding profile is included in new security policies by default:<br>a. Select **Policies > Security** and **Add** a new security policy.<br>b. Select the **Actions** tab and ensure the **Log Forwarding** field shows the *default* profile selected:<br> 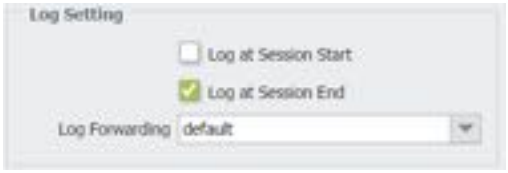 <br><br>5. Confirm the *default* log forwarding profile is included in new security zones by default:<br>a. Select **Network > Zones** and **Add** a new security zone.<br>b. Ensure the **Log Setting** field shows the *default* log forwarding profile selected:<br> <br> |
| • Override a default log forwarding profile. | If you have an existing default log forwarding profile, and you do not want the log forwarding settings defined in that profile to be applied to a new security policy or a new security zone, continue to modify the **Log Setting** field in the policy or zone according to your preference. |

# Monitor the Firewall Using SNMP

All Palo Alto Networks firewalls support standard networking SNMP management information base (MIB) modules as well as proprietary Enterprise MIB modules. You can configure an SNMP manager to get statistics from the firewall. For example, you could configure your SNMP manager to monitor the interfaces, active sessions, concurrent sessions, session utilization percentage, temperature, and/or system uptime on the firewall.

Palo Alto Networks firewalls support SNMP GET requests only; SNMP SET requests are not supported.

| **Set Up SNMP Monitoring** | | |
|---|---|---|
| Step 1 | Enable the interface to allow inbound SNMP requests. | • If you will be receiving SNMP GET messages on the MGT interface, select **Device > Setup > Management** and click to Edit the Management Interface Settings section of the screen. In the **Services** section, select the **SNMP** check box and then click **OK**.<br>• If you will be receiving SNMP GET messages on a different interface, you must associate a management profile with the interface and enable SNMP management. |
| Step 2 | From the web interface on the firewall, configure the settings to allow the SNMP agent on the firewall to respond to incoming GET requests from the SNMP manager. | 1. Select **Device > Setup > Operations > SNMP Setup**.<br>2. Specify the **Physical Location** of the firewall and the name or email address of an administrative **Contact**.<br>3. Select the SNMP **Version** and then enter the configuration details as follows (depending on which SNMP version you are using) and then click **OK**:<br>• **V2c**—Enter the **SNMP Community String** that will allow the SNMP manager access to the SNMP agent on the firewall. The default value is **public**, however because this is a well-known community string, it is a best practice to use a value that is not easily guessed.<br>• **V3**—You must create at least one View and one User in order to use SNMPv3. The view specifies which management information the manager has access to. If you want to allow access to all management information, just enter the top-level **OID** of .1.3.6.1 and specify the **Option** as **include** (you can also create views that exclude certain objects). Use **0xf0** as the **Mask**. Then when you create a user, select the **View** you just created and specify the **Auth Password** and **Priv Password**.<br>The authentication settings (the community string for V2c or the username and passwords for V3) configured on the firewall must match the value configured on the SNMP manager.<br>4. Click **OK** to save the settings.<br>5. Click **Commit** to save the SNMP settings. |
| Step 3 | Enable the SNMP manager to interpret firewall statistics. | Load the PAN-OS MIB files into your SNMP management software and, if necessary, compile them. Refer to the documentation for your SNMP manager for specific instructions on how to do this. |


© Palo Alto Networks, Inc.

| Set Up SNMP Monitoring (Continued) | | |
|---|---|---|
| Step 4 | Identify the statistics you want to monitor. | Using a MIB browser, walk the PAN-OS MIB files to identify the object identifiers (OIDs) that correspond to the statistics you want to monitor. For example, suppose you want to monitor Session Utilization Percentage on the firewall. Using a MIB browser you will see that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.1.2.3.1.0 in the PAN-COMMON-MIB. |
| Step 5 | Configure the SNMP management software to monitor the OIDs you are interested in. | Refer to the documentation for your SNMP manager for specific instructions on how to do this. |
| Step 6 | After you complete the configuration on both the firewall and the SNMP manager, you can begin monitoring the firewall from your SNMP management software. | The following is an example of how an SNMP manager displays real-time session utilization percentage statistics for a monitored PA-500 firewall: 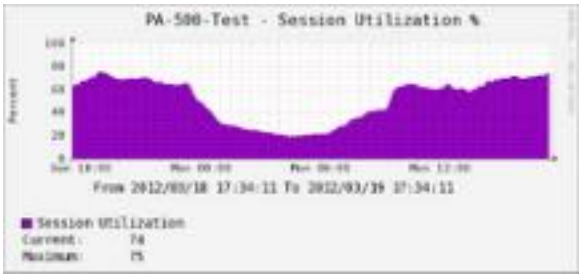  |

# Monitor the Firewall Using NetFlow

NetFlow is an industry-standard protocol that enables the firewall to record statistics on the IP traffic that traverses its interfaces. The firewall exports the statistics as NetFlow fields to a NetFlow collector. The NetFlow collector is a server you use to analyze network traffic for security, administration, accounting and troubleshooting purposes. All firewalls support NetFlow Version 9 except the PA-4000 Series and PA-7050 firewalls. The firewalls support only unidirectional NetFlow, not bidirectional. You can enable NetFlow exports on all interface types except HA, log card, or decrypt mirror. To identify firewall interfaces in a NetFlow collector, see Identify Firewall Interfaces in External Monitoring Systems.

The firewall supports standard and enterprise (PAN-OS specific) NetFlow templates. NetFlow collectors require templates to decipher the exported fields. The firewall selects a template based on the type of data it exports: IPv4 or IPv6 traffic, with or without NAT, and with standard or enterprise-specific fields. For a list of supported templates and field definitions, see NetFlow Templates. The firewall periodically refreshes templates to apply any changes to them. You configure the refresh frequency according to the requirements of the particular NetFlow collector you use.

| **Monitor the Firewall Using NetFlow** | | |
|---|---|---|
| Step 1 | Create a NetFlow server profile. | 1. Select **Device > Server Profiles > NetFlow** and click **Add**.<br>2. Enter a **Name** for the profile.<br>3. Specify the frequency at which the firewall refreshes templates (in **Minutes** or **Packets**) and exports records (**Active Timeout** in minutes).<br>4. Select the **PAN-OS Field Types** check box if you want the firewall to export App-ID and User-ID fields.<br>5. For each NetFlow collector (up to two per profile) that will receive fields, click **Add** and enter an identifying server **Name**, hostname or IP address (**NetFlow Server**), and access **Port** (default 2055).<br>6. Click **OK** to save the profile. |
| Step 2 | Assign the NetFlow server profile to the interfaces that carry the traffic you want to analyze.<br><br>In this example, you assign the profile to an existing Layer 3 interface. | 1. Select **Network > Interfaces > Ethernet** and click an **Interface** name to edit it.<br>2. In the **NetFlow Profile** drop-down, select the NetFlow server profile and click **OK**.<br>3. Click **Commit** and **OK**. |


© Palo Alto Networks, Inc.

# NetFlow Templates

The Palo Alto Networks firewall supports the following NetFlow templates:

| Template | ID |
|---|---|
| IPv4 Standard | 256 |
| IPv4 Enterprise | 257 |
| IPv6 Standard | 258 |
| IPv6 Enterprise | 259 |
| IPv4 with NAT Standard | 260 |
| IPv4 with NAT Enterprise | 261 |

The following table lists the NetFlow fields that the firewall can send, along with the templates that define them:

| Value | Field | Description | Templates |
|---|---|---|---|
| 1 | IN_BYTES | Incoming counter with length N * 8 bits for the number of bytes associated with an IP flow. By default, N is 4. | All templates |
| 2 | IN_PKTS | Incoming counter with length N * 8 bits for the number of packets associated with an IP glow. By default, N is 4 | All templates |
| 4 | PROTOCOL | IP protocol byte. | All templates |
| 5 | TOS | Type of Service byte setting when entering the incoming interface. | All templates |
| 6 | TCP_FLAGS | Total of all the TCP flags in this flow. | All templates |
| 7 | L4_SRC_PORT | TCP/UDP source port number (for example, FTP, Telnet, or equivalent). | All templates |
| 8 | IPV4_SRC_ADDR | IPv4 source address. | IPv4 standard<br>IPv4 enterprise<br>IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 10 | INPUT_SNMP | Input interface index. The value length is 2 bytes by default, but higher values are possible. For details on how Palo Alto Networks firewalls generate interface indexes, see Identify Firewall Interfaces in External Monitoring Systems. | All templates |
| 11 | L4_DST_PORT | TCP/UDP destination port number (for example, FTP, Telnet, or equivalent). | All templates |

| Value | Field | Description | Templates |
|---|---|---|---|
| 12 | IPV4_DST_ADDR | IPv4 destination address. | IPv4 standard<br>IPv4 enterprise<br>IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 14 | OUTPUT_SNMP | Output interface index. The value length is 2 bytes by default, but higher values are possible. For details on how Palo Alto Networks firewalls generate interface indexes, see Identify Firewall Interfaces in External Monitoring Systems. | All templates |
| 21 | LAST_SWITCHED | System uptime in milliseconds when the last packet of this flow was switched. | All templates |
| 22 | FIRST_SWITCHED | System uptime in milliseconds when the first packet of this flow was switched. | All templates |
| 27 | IPV6_SRC_ADDR | IPv6 source address. | IPv6 standard<br>IPv6 enterprise |
| 28 | IPV6_DST_ADDR | IPv6 destination address. | IPv6 standard<br>IPv6 enterprise |
| 32 | ICMP_TYPE | Internet Control Message Protocol (ICMP) packet type. This is reported as:<br>ICMP Type * 256 + ICMP code | All templates |
| 61 | DIRECTION | Flow direction:<br>• 0 = ingress<br>• 1 = egress | All templates |
| 148 | flowId | An identifier of a flow that is unique within an *observation domain*. You can use this *information element* to distinguish between different flows if *flow keys* such as IP addresses and port numbers are not reported or are reported in separate records. The flowID corresponds to the session ID field in Traffic and Threat logs. | All templates |
| 233 | firewallEvent | Indicates a firewall event:<br>• 0 = Ignore (invalid)<br>• 1 = Flow created<br>• 2 = Flow deleted<br>• 3 = Flow denied<br>• 4 = Flow alert<br>• 5 = Flow update (the session state changed from active to deny) | All templates |

 © Palo Alto Networks, Inc.

| Value | Field | Description | Templates |
|---|---|---|---|
| 225 | postNATSourceIPv4Address | The definition of this information element is identical to that of sourceIPv4Address, except that it reports a modified value that a NAT middlebox function caused after the packet passed the *observation point*. | IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 226 | postNATDestinationIPv4Address | The definition of this information element is identical to that of destinationIPv4Address, except that it reports a modified value that a NAT middlebox function caused after the packet passed the observation point. | IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 227 | postNAPTSourceTransportPort | The definition of this information element is identical to that of sourceTransportPort, except that it reports a modified value that a Network Address Port Translation (NAPT) middlebox function caused after the packet passed the observation point. | IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 228 | postNAPTDestinationTransportPort | The definition of this information element is identical to that of destinationTransportPort, except that it reports a modified value that a Network Address Port Translation (NAPT) middlebox function caused after the packet passed the observation point. | IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 346 | privateEnterpriseNumber | This is a unique private enterprise number that identifies Palo Alto Networks: 25461. | IPv4 enterprise<br>IPv4 with NAT enterprise<br>IPv6 enterprise |
| 56701 | App-ID | The name of an application that App-ID identified. The name can be up to 32 bytes. | IPv4 enterprise<br>IPv4 with NAT enterprise<br>IPv6 enterprise |
| 56702 | User-ID | A username that User-ID identified. The name can be up to 64 bytes. | IPv4 enterprise<br>IPv4 with NAT enterprise<br>IPv6 enterprise |

# Identify Firewall Interfaces in External Monitoring Systems

When you use a NetFlow collector (see Monitor the Firewall Using NetFlow) or SNMP manager (see Monitor the Firewall Using SNMP) to monitor traffic flows, an interface index (SNMP ifindex object) identifies the firewall interface that carried a particular flow. The formula that the Palo Alto Networks firewall uses to calculate interface indexes varies by platform and whether the interface is physical or logical.

You cannot use SNMP to monitor logical interfaces, only physical interfaces. You can use NetFlow to monitor logical or physical interfaces.

Most NetFlow collectors use SNMP to determine the name of a physical interface based on the SNMP interface index.

Physical interface indexes have a range of 1-9999, which the firewall calculates as follows:

| Firewall Platform | Calculation | Example Interface Index |
|---|---|---|
| Non-chassis based:<br>VM-Series, PA-200, PA-500, PA-2000 Series, PA-3000 Series, PA-4000 Series, PA-5000 Series<br>The PA-4000 Series platform supports SNMP but not NetFlow. | MGT port + physical port offset<br>• MGT port—This is a constant that depends on the platform:<br>  • 2 for hardware-based firewalls (for example, the PA-5000 Series firewall)<br>  • 1 for the VM-Series firewall<br>• Physical port offset—This is the physical port number. | PA-5000 Series firewall, Eth1/4 =<br>2 (MGT port) + 4 (physical port) = **6** |
| Chassis based:<br>PA-7050 firewalls<br>This platform supports SNMP but not NetFlow. | (Max. ports * slot) + physical port offset + MGT port<br>• Maximum ports—This is a constant of 64.<br>• Slot—This is the chassis slot number of the network interface card.<br>• Physical port offset—This is the physical port number.<br>• MGT port—This is a constant of 5 for PA-7050 firewalls. | PA-7050 firewall, Eth3/9 =<br>[64 (max. ports) * 3 (slot)] + 9 (phyical port) + 5 (MGT port) = **206** |

Logical interface indexes for all platforms are nine-digit numbers that the firewall calculates as follows:

| Interface Type | Range | Digit 9 | Digits 7-8 | Digits 5-6 | Digits 1-4 | Example Interface Index |
|---|---|---|---|---|---|---|
| Layer 3 subinterface | 101010001-199999999 | Type: 1 | Interface slot: 1-9 (01-09) | Interface port: 1-9 (01-09) | Subinterface: suffix 1-9999 (0001-9999) | Eth1/5.22 = 100000000 (type) + 100000 (slot) + 50000 (port) + 22 (suffix) = **101050022** |
| Layer 2 subinterface | 101010001-199999999 | Type: 1 | Interface slot: 1-9 (01-09) | Interface port: 1-9 (01-09) | Subinterface: suffix 1-9999 (0001-9999) | Eth2/3.6 = 100000000 (type) + 200000 (slot) + 30000 (port) + 6 (suffix) = **102030006** |

 © Palo Alto Networks, Inc.

| Interface Type | Range | Digit 9 | Digits 7-8 | Digits 5-6 | Digits 1-4 | Example Interface Index |
|---|---|---|---|---|---|---|
| Vwire subinterface | 101010001-199999999 | Type: 1 | Interface slot: 1-9 (01-09) | Interface port: 1-9 (01-09) | Subinterface: suffix 1-9999 (0001-9999) | Eth4/2.312 = 100000000 (type) + 400000 (slot) + 20000 (port) + 312 (suffix) = **104020312** |
| VLAN | 200000001-200009999 | Type: 2 | 00 | 00 | VLAN suffix: 1-9999 (0001-9999) | VLAN.55 = 200000000 (type) + 55 (suffix) = **200000055** |
| Loopback | 300000001-300009999 | Type: 3 | 00 | 00 | Loopback suffix: 1-9999 (0001-9999) | Loopback.55 = 300000000 (type) + 55 (suffix) = **300000055** |
| Tunnel | 400000001-400009999 | Type: 4 | 00 | 00 | Tunnel suffix: 1-9999 (0001-9999) | Tunnel.55 = 400000000 (type) + 55 (suffix) = **400000055** |
| Aggregate group | 500010001-500089999 | Type: 5 | 00 | AE suffix: 1-8 (01-08) | Subinterface: suffix 1-9999 (0001-9999) | AE5.99 = 500000000 (type) + 50000 (AE Suffix) + 99 (suffix) = **500050099** |

# Manage Reporting

The reporting capabilities on the firewall allow you to keep a pulse on your network, validate your policies, and focus your efforts on maintaining network security for keeping your users safe and productive.


- About Reports
- View Reports
- Disable Predefined Reports
- Generate Custom Reports
- Generate Botnet Reports
- Manage PDF Summary Reports
- Generate User/Group Activity Reports
- Manage Report Groups
- Schedule Reports for Email Delivery


 © Palo Alto Networks, Inc.

## About Reports

The firewall includes predefined reports that you can use as-is, or you can build custom reports that meet your needs for specific data and actionable tasks, or you can combine predefined and custom reports to compile information you need. The firewall provides the following types of reports:

- **Predefined Reports—**Allow you to view a quick summary of the traffic on your network. A suite of predefined reports are available in four categories—Applications, Traffic, Threat, and URL Filtering. See View Reports.
- **User or Group Activity Reports—**Allow you to schedule or create an on-demand report on the application use and URL activity for a specific user or for a user group. The report includes the URL categories and an estimated browse time calculation for individual users. See Generate User/Group Activity Reports.
- **Custom Reports**—Create and schedule custom reports that show exactly the information you want to see by filtering on conditions and columns to include. You can also include query builders for more specific drill down on report data. See Generate Custom Reports.
- **PDF Summary Reports**—Aggregate up to 18 predefined or custom reports/graphs from Threat, Application, Trend, Traffic, and URL Filtering categories into one PDF document. See Manage PDF Summary Reports.
- **Botnet Reports**—Allow you to use behavior-based mechanisms to identify potential botnet-infected hosts in the network. See Interpret Botnet Report Output.
- **Report Groups**—Combine custom and predefined reports into report groups and compile a single PDF that is emailed to one or more recipients. See Manage Report Groups.

Reports can be generated on demand, on a recurring schedule, and can be scheduled for email delivery.

# View Reports

The firewall provides an assortment of over 40 predefined reports that are generated every day; these reports can be viewed directly on the firewall. In addition to these reports, you can view *scheduled* custom reports and summary reports.

About 200 MB of storage is allocated for saving reports on the firewall. This storage space is not user-configurable, and older reports are purged to store recent reports. Therefore, for long-term retention of reports, you can either export the reports or schedule the reports for email delivery. To disable selected reports and conserve system resources on the firewall, see Disable Predefined Reports.

User/group activity reports must be generated on demand or scheduled for email delivery. Unlike the other reports, these reports cannot be saved on the firewall.

**View Reports**

Step 1 Select **Monitor > Reports**.

The reports are chunked into sections on the right-hand side of the window: **Custom Reports**, **Application Reports**, **Traffic Reports**, **Threat Reports**, **URL Filtering Reports**, and **PDF Summary Reports**.

Step 2 Select a report to view. When you select a report, the previous day's report is displayed on screen.

To view reports for any of the previous days, select an available date from the calendar at the bottom of the page and select a report within the same section. If you change sections, the time selection is reset.

Step 3 To view a report offline, you can export the report to PDF, CSV or to XML formats. Click **Export to PDF**, **Export to CSV**, or **Export to XML** at the bottom of the page. Then print or save the file.

 © Palo Alto Networks, Inc.

## Disable Predefined Reports

The firewall includes about 40 predefined reports that are automatically generated each day. If you do not use some or all of these predefined reports, you can disable selected reports and conserve system resources on the firewall.

Before disabling one or more predefined reports, make sure that the report is not included in a Group Report or a PDF Report. If the disabled predefined report is included in a group or PDF report, the Group/PDF report will be rendered without any data.

**Disable Predefined Reports**

1. Select **Device > Setup > Management** on the firewall.
2. Click the Edit icon in the Logging and Reporting Settings section and select the **Log Export and Reporting** tab.
3. To disable reports:
   - Clear the check box corresponding to each report that you want to disable.
   - Select **Deselect All** to disable all predefined reports.

4. Click **OK**, and **Commit** the changes.

## Generate Custom Reports

In order to create purposeful custom reports, you must consider the attributes or key pieces of information that you want to retrieve and analyze. This consideration guides you in making the following selections in a custom report:

| Selection | Description |
|---|---|
| Data Source | The data file that is used to generate the report. The firewall offers two types of data sources—Summary databases and Detailed logs.<br>• Summary databases are available for traffic, threat, and application statistics. The firewall aggregates the detailed logs on traffic, application, and threat at 15-minute intervals. The data is condensed—duplicate sessions are grouped together and incremented with a repeat counter, and some attributes (or columns) are not included in the summary—to allow faster response time when generating reports.<br>• Detailed logs are itemized and are a complete listing of all the attributes (or columns) that pertain to the log entry. Reports based on detailed logs take much longer to run and are not recommended unless absolutely necessary. |
| Attributes | The columns that you want to use as the match criteria. The attributes are the columns that are available for selection in a report. From the list of **Available Columns**, you can add the selection criteria for matching data and for aggregating the details (the **Selected Columns**). |
| Sort By/ Group By | The **Sort By** and the **Group By** criteria allow you to organize/segment the data in the report; the sorting and grouping attributes available vary based on the selected data source.<br>The Sort By option specifies the attribute that is used for aggregation. If you do not select an attribute to sort by, the report will return the first N number of results without any aggregation.<br>The Group By option allows you to select an attribute and use it as an anchor for grouping data; all the data in the report is then presented in a set of top 5, 10, 25 or 50 groups. For example, when you select Hour as the Group By selection and want the top 25 groups for a 24-hr time period, the results of the report will be generated on an hourly basis over a 24-hr period. The first column in the report will be the hour and the next set of columns will be the rest of your selected report columns. |

 © Palo Alto Networks, Inc.

| Selection | Description |
|---|---|
| | The following example illustrates how the **Selected Columns** and **Sort By**/**Group By** criteria work together when generating reports:<br> 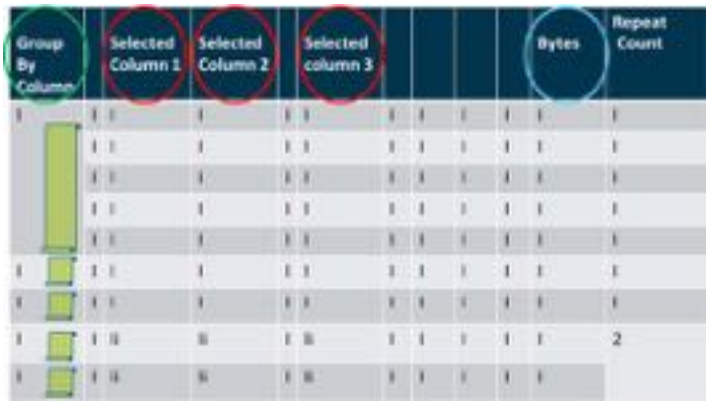 <br>The columns circled in red (above) depict the columns selected, which are the attributes that you match against for generating the report. Each log entry from the data source is parsed and these columns are matched on. If multiple sessions have the same values for the selected columns, the sessions are aggregated and the repeat count (or sessions) is incremented.<br>The column circled in blue indicates the chosen sort order. When the sort order (**Sort By**) is specified, the data is sorted (and aggregated) by the selected attribute.<br>The column circled in green indicates the **Group By** selection, which serves as an anchor for the report. The **Group By** column is used as a match criteria to filter for the top N groups. Then, for each of the top N groups, the report enumerates the values for all the other selected columns. |

| Selection | Description |
|---|---|

For example, if a report has the following selections
:

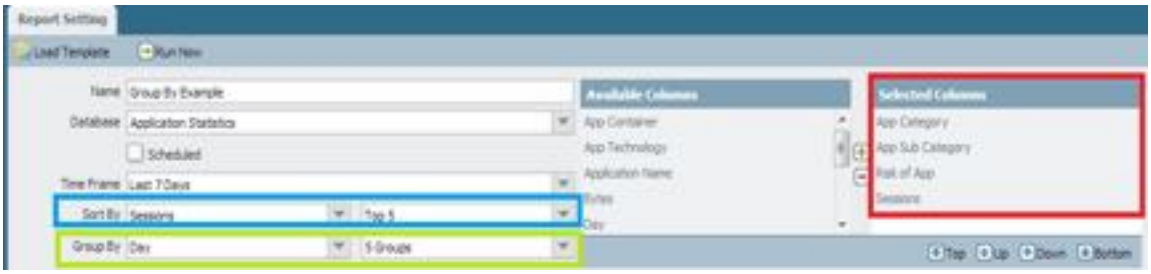


The output will display as follows:

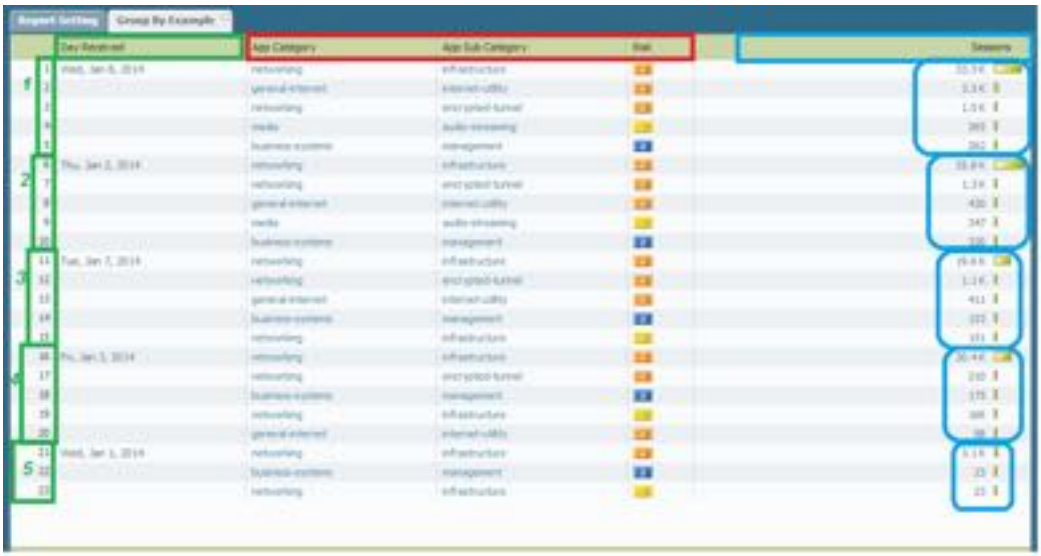

The report is anchored by **Day** and sorted by **Sessions**. It lists the 5 days (**5 Groups**) with maximum traffic in the **Last 7 Days** time frame. The data is enumerated by the **Top 5** sessions for each day for the selected columns—**App Category**, **App Subcategory** and **Risk**.

| | |
|---|---|
| Time Period | The date range for which you want to analyze data. You can define a custom range or select a time period ranging from last 15 minutes to the last 30 days. The reports can be run on demand or scheduled to run at a daily or weekly cadence. |
| Query Builder | The query builder allows you to define specific queries to further refine the selected attributes. It allows you see just what you want in your report using **and** and **or** operators and a match criteria, and then include or exclude data that matches or negates the query in the report. Queries enable you to generate a more focused collation of information in a report. |

 © Palo Alto Networks, Inc.

## Generate Custom Reports

1. Select **Monitor > Manage Custom Reports**.
2. Click **Add** and then enter a **Name** for the report.

   To base a report on an predefined template, click **Load Template** and choose the template. You can then edit the template and save it as a custom report.
3. Select the database to use for the report.

Each time you create a custom report, a **Log View** report is automatically created. This report show the logs that were used to build the custom report. The log view report uses the same name as the custom report, but appends the phrase (Log View) to the report name.

When creating a report group, you can include the log view report with the custom report. For more information, see Manage Report Groups.

4. Select the **Scheduled** check box to run the report each night. The report is then available for viewing in the **Reports** column on the side.
5. Define the filtering criteria. Select the **Time Frame**, the **Sort By** order, **Group By** preference, and select the columns that must display in the report.
6. (Optional) Select the **Query Builder** attributes if you want to further refine the selection criteria. To build a report query, specify the following and click **Add**. Repeat as needed to construct the full query.
   - **Connector**—Choose the connector (and/or) to precede the expression you are adding.
   - **Negate**—Select the check box to interpret the query as a negation. If, for example, you choose to match entries in the last 24 hours and/or are originating from the untrust zone, the negate option causes a match on entries that are not in the past 24 hours and/or are not from the untrust zone.
   - **Attribute**—Choose a data element. The available options depend on the choice of database.
   - **Operator**—Choose the criterion to determine whether the attribute applies (such as =). The available options depend on the choice of database.
   - **Value**—Specify the attribute value to match.

   For example, the following figure (based on the Traffic Log database) shows a query that matches if the traffic log entry was received in the past 24 hours and is from the "untrust" zone.

7. To test the report settings, select **Run Now**. Modify the settings as required to change the information that is displayed in the report.
8. Click **OK** to save the custom report.

**Generate Custom Reports (Continued)**

**Examples of Custom Reports**

If you want to set up a simple report in which you use the traffic summary database from the last 30 days, and sort the data by the top 10 sessions and these sessions are grouped into 5 groups by day of the week. You would set up the custom report to look like this:

And the PDF output for the report would look as follows:

 © Palo Alto Networks, Inc.

**Generate Custom Reports (Continued)**

Now, if you want to use the query builder to generate a custom report that represents the top consumers of network resources within a user group, you would set up the report to look like this:

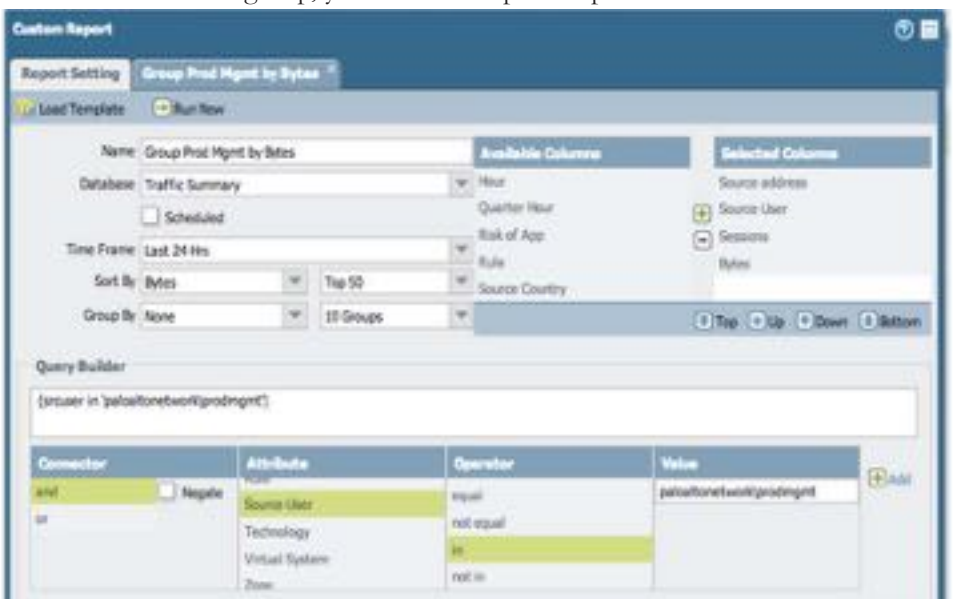


The report would display the top users in the product management user group sorted by bytes, as follows:

# Generate Botnet Reports

The botnet report enables you to use heuristic and behavior-based mechanisms to identify potential malware- or botnet-infected hosts in your network. To evaluate botnet activity and infected hosts, the firewall correlates user and network activity data in Threat, URL, and Data Filtering logs with the list of malware URLs in PAN-DB, known dynamic DNS domain providers, and domains registered within the last 30 days. You can configure the report to identify hosts that visited those sites, as well as hosts that communicated with Internet Relay Chat (IRC) servers or that used unknown applications. Malware often use dynamic DNS to avoid IP blacklisting, while IRC servers often use bots for automated functions.

The firewall requires Threat Prevention and URL Filtering licenses to use the botnet report.

▲ Configure a Botnet Report

▲ Interpret Botnet Report Output

## Configure a Botnet Report

You can schedule a botnet report or run it on demand. The firewall generates scheduled botnet reports every 24 hours because behavior-based detection requires correlating traffic across multiple logs over that timeframe.

| **Configure a Botnet Report** | |
|---|---|
| Step 1 Define the types of traffic that indicate possible botnet activity. | 1. Select **Monitor > Botnet** and click **Configuration** on the right side of the page.<br>2. **Enable** and define the **Count** for each type of HTTP Traffic that the report will include.<br>The **Count** values represent the minimum number of events of each traffic type that must occur for the report to list the associated host with a higher confidence score (higher likelihood of botnet infection). If the number of events is less than the **Count**, the report will display a lower confidence score or (for certain traffic types) won't display an entry for the host. For example, if you set the **Count** to three for **Malware URL visit**, then hosts that visit three or more known malware URLs will have higher scores than hosts that visit less than three. For details, see Interpret Botnet Report Output.<br>3. Define the thresholds that determine whether the report will include hosts associated with traffic involving Unknown TCP or Unknown UDP applications.<br>4. Select the **IRC** check box to include traffic involving IRC servers.<br>5. Click **OK** to save the report configuration. |

 © Palo Alto Networks, Inc.

| Configure a Botnet Report (Continued) | |
|---|---|
| Step 2 Schedule the report or run it on demand. | 1. Click **Report Setting** on the right side of the page.<br>2. Select a time interval for the report in the **Test Run Time Frame** drop-down.<br>3. Select the **No. of Rows** to include in the report.<br>4. (Optional) **Add** queries to the Query Builder to filter the report output by attributes such as source/destination IP addresses, users, or zones.<br>For example, if you know in advance that traffic initiated from the IP address 10.3.3.15 contains no potential botnet activity, you can add `not (addr.src in 10.0.1.35)` as a query to exclude that host from the report output. For details, see Interpret Botnet Report Output.<br>5. Select **Scheduled** to run the report daily or click **Run Now** to run the report immediately.<br>6. Click **OK** and **Commit**. |

## Interpret Botnet Report Output

The botnet report displays a line for each host that is associated with traffic you defined as suspicious when configuring the report. For each host, the report displays a confidence score of 1 to 5 to indicate the likelihood of botnet infection, where 5 indicates the highest likelihood. The scores correspond to threat severity levels: 1 is informational, 2 is low, 3 is medium, 4 is high, and 5 is critical. The firewall bases the scores on:

- **Traffic type**—Certain HTTP traffic types are more likely to involve botnet activity. For example, the report assigns a higher confidence to hosts that visit known malware URLs than to hosts that browse to IP domains instead of URLs, assuming you defined both those activities as suspicious.
- **Number of events**—Hosts that are associated with a higher number of suspicious events will have higher confidence scores based on the thresholds (**Count** values) you define when you Configure a Botnet Report.
- **Executable downloads**—The report assigns a higher confidence to hosts that download executable files. Executable files are a part of many infections and, when combined with the other types of suspicious traffic, can help you prioritize your investigations of compromised hosts.

When reviewing the report output, you might find that the sources the firewall uses to evaluate botnet activity (for example, the list of malware URLs in PAN-DB) have gaps. You might also find that these sources identify traffic that you consider safe. To compensate in both cases, you can add query filters when you Configure a Botnet Report.

## Manage PDF Summary Reports

PDF summary reports contain information compiled from existing reports, based on data for the top 5 in each category (instead of top 50). They also contain trend charts that are not available in other reports.

| **Generate PDF Summary Reports** | |
|---|---|
| Step 1 Set up a **PDF Summary Report**. | 1. Select **Monitor > PDF Reports > Manage PDF Summary**.<br>2. Click **Add** and then enter a **Name** for the report.<br>3. Use the drop-down for each report group and select one or more of the elements to design the PDF Summary Report. You can include a maximum of 18 report elements.<br> 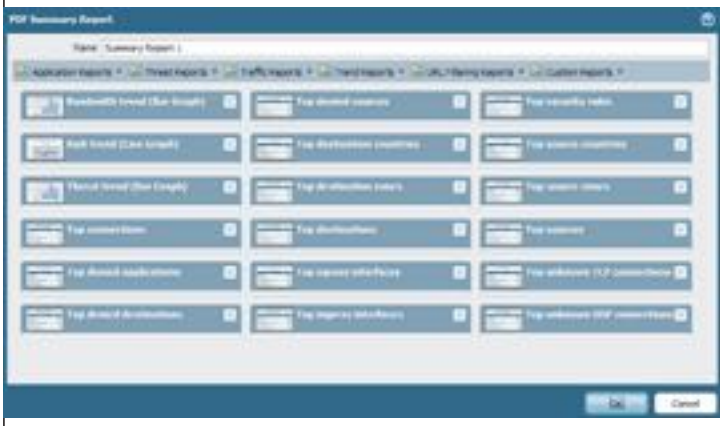 <br>• To remove an element from the report, click the **x** icon or clear the selection from the drop-down for the appropriate report group.<br>• To rearrange the reports, drag and drop the icons to another area of the report.<br>4. Click **OK t**o save the report.<br>5. **Commit** the changes. |


© Palo Alto Networks, Inc.

| Generate PDF Summary Reports | |
|---|---|
| Step 2 View the report. | To download and view the PDF Summary Report, see View Reports. 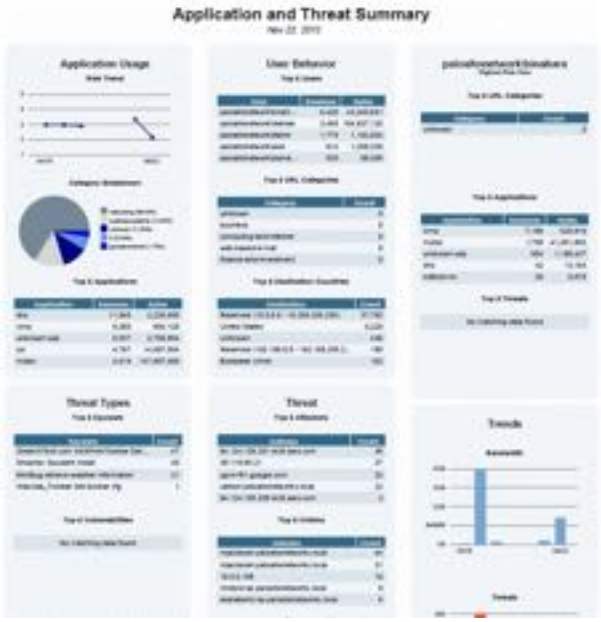  |

## Generate User/Group Activity Reports

User/Group activity reports summarize the web activity of individual users or user groups. Both reports include the same information with a couple exceptions—**Browsing Summary by URL Category** and **Browse time calculations** are included in User Activity Reports, but are not included in Group Activity Reports.

The User Activity Report includes two variables—the **Average Browse Time (sec)** and the **Page Load Threshold (sec) i**n the Logging and Reporting sections of the **Device > Setup > Management tab**—when calculating the browse time. The **Average Browse Time** is the average time that the admin estimates that a user will browse a web page. Any request made after the average browse time has elapsed will be considered a new browsing activity. The calculation uses Container Pages (logged in the URL filtering logs) as the basis and ignores any new web pages that are loaded between the time of the first request (start time) and the average browse time. Example: If the average browse time setting is 2 minutes and a user opens a web page and views that page for 5 minutes, the browse time for that page will still be 2 minutes. This is done because there is no way to determine how long a user views a given page.The Average Browse Time calculation ignores sites categorized as web advertisements and content delivery networks. The **Page Load Threshold** is the estimated time for page elements to load on the page. Any request that occurs between the first page load and the page load threshold is assumed to be elements of the page. Any requests that occur outside of the page load threshold are assumed to be the user clicking a link within the page.

User-ID must be configured on the firewall in order to access the list of user/user groups.

### Generate User/Group Activity Reports

1. Select **Monitor > PDF Reports > User Activity Report**.
2. Click **Add** and then enter a **Name** for the report.
3. Create the report:
   - For a User Activity Report: Select **User** and enter the **Username** or **IP address** (IPv4 or IPv6) of the user who will be the subject of the report.
   - For Group Activity Report: Select **Group** and select the **Group Name** for which to retrieve user group information in the report.
4. Select the time frame for the report from the drop-down.

   The number of logs that are analyzed in a user activity report is determined by the number of rows defined on the **Max Rows in User Activity Report** on the Logging and Reporting Settings section on **Device >Setup > Management**.
5. Select **Include Detailed Browsing** to include detailed URL logs in the report.

   The detailed browsing information can include a large volume of logs (thousands of logs) for the selected user or user group and can make the report very large.
6. To run the report on demand, click **Run Now**.
7. To save the report, click **OK**. User/Group activity reports cannot be saved on the firewall; to schedule the report for email delivery, see Schedule Reports for Email Delivery.

 © Palo Alto Networks, Inc.

# Manage Report Groups

Report groups allow you to create sets of reports that the system can compile and send as a single aggregate PDF report with an optional title page and all the constituent reports included.

| Set up Report Groups | |
|---|---|
| Step 1 Set up report groups.<br><br>You must set up a **Report Group** to email report(s). | 1. Create a Server Profile for your email server.<br>2. Define the **Report Group**. A report group can compile predefined reports, PDF Summary reports, custom reports, and Log View report into a single PDF.<br>a. Select **Monitor > Report Group**.<br>b. Click **Add** and then enter a **Name** for the report group.<br>c. (Optional) Select **Title Page** and add a **Title** for the PDF output.<br>d. Select reports from the left column and click **Add** to move each report to the report group on the right.<br><br> 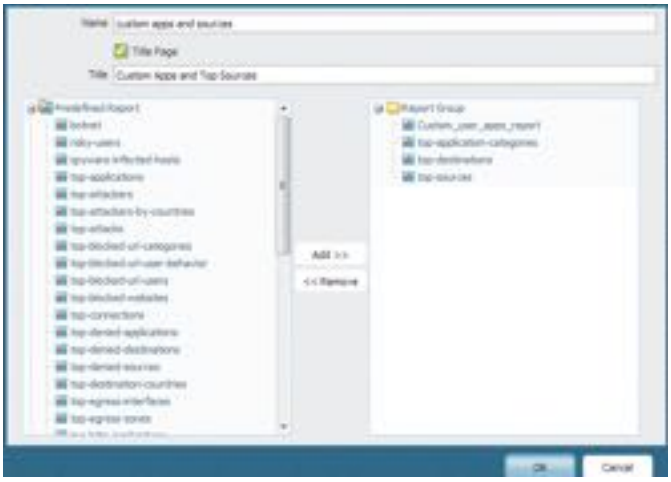 <br><br>The **Log View** report is a report type that is automatically created each time you create a custom report and uses the same name as the custom report. This report will show the logs that were used to build the contents of the custom report.<br><br>To include the log view data, when creating a report group, add your custom report under the **Custom Reports** list and then add the log view report by selecting the matching report name from the **Log View** list. The report will include the custom report data and the log data that was used to create the custom report.<br>e. Click **OK** to save the settings.<br>f. To use the report group, see Schedule Reports for Email Delivery. |

## Schedule Reports for Email Delivery

Reports can be scheduled for daily delivery or delivered weekly on a specified day. Scheduled reports are executed starting at 2:00 AM, and email delivery starts after all scheduled reports have been generated.

**Schedule Reports for Email Delivery**

1. Select **Monitor > PDF Reports > Email Scheduler**.
2. Select the **Report Group** for email delivery. To set up a report group; see Manage Report Groups.
3. Select the frequency at which to generate and send the report in **Recurrence**.
4. Select the email server profile to use for delivering the reports. To set up an email server profile, see Create a Server Profile for your email server.
5. The **Override Recipient email(s)** allows you to send this report exclusively to the recipients specified in this field. When you add recipients to the **Override Recipient email(s)**, the report is not sent to the recipients configured in the email server profile. Use this option for those occasions when the report is for the attention of someone other than the administrators or recipients defined in the email server profile.

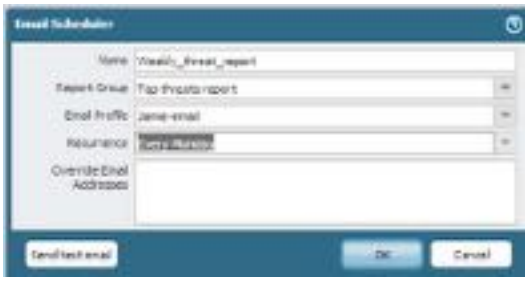

 © Palo Alto Networks, Inc.

# Syslog Field Descriptions

This is a list of the standard fields for each of the five log types that are forwarded to an external server. For ease of parsing, the comma is the delimiter; each field is a comma-separated value (CSV) string. The FUTURE_USE tag applies to fields that the devices do not currently implement.

WildFire logs are a subtype of threat logs and use the same Syslog format.


- Traffic Logs
- Threat Logs
- HIP Match Logs
- Config Logs
- System Logs
- Syslog Severity
- Custom Log/Event Format
- Escape Sequences


## Traffic Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Ingress Interface, Egress Interface, Log Forwarding Profile, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Bytes, Bytes Sent, Bytes Received, Packets, Start Time, Elapsed Time, Category, FUTURE_USE, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Packets Sent, Packets Received, Session End Reason*

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Specifies type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of traffic log are as follows:<br>• start—session started<br>• end—session ended<br>• drop—session dropped before the application is identified and there is no rule that allows the session.<br>• deny—session dropped after the application is identified and there is a rule to block or no rule that allows the session. |
| Generated Time (time_generated) | Time the log was generated on the dataplane |

| Field Name | Description |
|---|---|
| Source IP (src) | Original session source IP address |
| Destination IP (dst) | Original session destination IP address |
| NAT Source IP (natsrc) | If Source NAT performed, the post-NAT Source IP address |
| NAT Destination IP (natdst) | If Destination NAT performed, the post-NAT Destination IP address |
| Rule Name (rule) | Name of the rule that the session matched |
| Source User (srcuser) | Username of the user who initiated the session |
| Destination User (dstuser) | Username of the user to which the session was destined |
| Application (app) | Application associated with the session |
| Virtual System (vsys) | Virtual System associated with the session |
| Source Zone (from) | Zone the session was sourced from |
| Destination Zone (to) | Zone the session was destined to |
| Ingress Interface (inbound_if) | Interface that the session was sourced form |
| Egress Interface (outbound_if) | Interface that the session was destined to |
| Log Forwarding Profile (logset) | Log Forwarding Profile that was applied to the session |
| Session ID (sessionid) | An internal numerical identifier applied to each session |
| Repeat Count (repeatcnt) | Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only |
| Source Port (sport) | Source port utilized by the session |
| Destination Port (dport) | Destination port utilized by the session |
| NAT Source Port (natsport) | Post-NAT source port |
| NAT Destination Port (natdport) | Post-NAT destination port |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Flags (flags) | 32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value:<br>• 0x80000000—session has a packet capture (PCAP)<br>• 0x02000000—IPv6 session<br>• 0x01000000—SSL session was decrypted (SSL Proxy)<br>• 0x00800000—session was denied via URL filtering<br>• 0x00400000—session has a NAT translation performed (NAT)<br>• 0x00200000—user information for the session was captured via the captive portal (Captive Portal)<br>• 0x00080000—X-Forwarded-For value from a proxy is in the source user field<br>• 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction)<br>• 0x00008000—session is a container page access (Container Page)<br>• 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above.<br>• 0x00000800—symmetric return was used to forward traffic for this session |
| Protocol (proto) | IP protocol associated with the session |
| Action (action) | Action taken for the session; values are allow or deny:<br>• Allow—session was allowed by policy<br>• Deny—session was denied by policy |
| Bytes (bytes) | Number of total bytes (transmit and receive) for the session |
| Bytes Sent (bytes_sent) | Number of bytes in the client-to-server direction of the session<br>Available on all models except the PA-4000 Series |
| Bytes Received (bytes_received) | Number of bytes in the server-to-client direction of the session<br>Available on all models except the PA-4000 Series |
| Packets (packets) | Number of total packets (transmit and receive) for the session |
| Start Time (start) | Time of session start |
| Elapsed Time (elapsed) | Elapsed time of the session |
| Category (category) | URL category associated with the session (if applicable) |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama |
| Source Location (srcloc) | Source country or Internal region for private addresses; maximum length is 32 bytes |

| Field Name | Description |
|---|---|
| Destination Location (dstloc) | Destination country or Internal region for private addresses. Maximum length is 32 bytes |
| Packets Sent (pkts_sent) | Number of client-to-server packets for the session<br>Available on all models except the PA-4000 Series |
| Packets Received (pkts_received) | Number of server-to-client packets for the session<br>Available on all models except the PA-4000 Series |
| Session End Reason (session_end_reason)<br>New in v6.1! | The reason a session terminated. If the termination had multiple causes, this field displays only the highest priority reason. The possible session end reason values are as follows, in order of priority (where the first is highest):<br>• threat—The firewall detected a threat associated with a reset, drop, or block (IP address) action.<br>• policy-deny—The session matched a security policy with a deny or drop action.<br>• tcp-rst-from-client—The client sent a TCP reset to the server.<br>• tcp-rst-from-server—The server sent a TCP reset to the client.<br>• resources-unavailable—The session dropped because of a system resource limitation. For example, the session could have exceeded the number of out-of-order packets allowed per flow or the global out-of-order packet queue.<br>• tcp-fin—One host or both hosts in the connection sent a TCP FIN message to close the session.<br>• tcp-reuse—A session is reused and the firewall closes the previous session.<br>• decoder—The decoder detects a new connection within the protocol (such as HTTP-Proxy) and ends the previous connection.<br>• aged-out—The session aged out.<br>• unknown—This value applies in the following situations:<br>  • Session terminations that the preceding reasons do not cover (for example, a `clear session all` command).<br>  • For logs generated in a PAN-OS release that does not support the session end reason field (releases older than PAN-OS 6.1), the value will be *unknown* after an upgrade to the current PAN-OS release or after the logs are loaded onto the firewall.<br>  • In Panorama, logs received from firewalls for which the PAN-OS version does not support session end reasons will have a value of *unknown*.<br>• n/a—This value applies when the traffic log type is not *end*. |

 © Palo Alto Networks, Inc.

## Threat Logs

**Format**: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Ingress Interface, Egress Interface, Log Forwarding Profile, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Miscellaneous, Threat ID, Category, Severity, Direction, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Content Type, PCAP_id, Filedigest, Cloud, FUTURE_USE, User Agent*, File Type*, X-Forwarded-For*, Referer*, Sender*, Subject*, Recipient*, Report ID*

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Specifies type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of threat log. Values include the following:<br>• data—Data pattern matching a Data Filtering profile.<br>• file—File type matching a File Blocking profile.<br>• flood—Flood detected via a Zone Protection profile.<br>• packet—Packet-based attack protection triggered by a Zone Protection profile.<br>• scan—Scan detected via a Zone Protection profile.<br>• spyware—Spyware detected via an Anti-Spyware profile.<br>• url—URL filtering log.<br>• virus—Virus detected via an Antivirus profile.<br>• vulnerability—Vulnerability exploit detected via a Vulnerability Protection profile.<br>• wildfire—A WildFire verdict generated when the firewall submits a file to WildFire per a WildFire Analysis profile and a verdict (malicious, grayware, or benign, depending on what you are logging) is logged in the WildFire Submissions log.<br>• wildfire-virus—Virus detected via an Antivirus profile. |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Source IP (src) | Original session source IP address |
| Destination IP (dst) | Original session destination IP address |
| NAT Source IP (natsrc) | If source NAT performed, the post-NAT source IP address |
| NAT Destination IP (natdst) | If destination NAT performed, the post-NAT destination IP address |
| Rule Name (rule) | Name of the rule that the session matched |
| Source User (srcuser) | Username of the user who initiated the session |
| Destination User (dstuser) | Username of the user to which the session was destined |

| Field Name | Description |
|---|---|
| Application (app) | Application associated with the session |
| Virtual System (vsys) | Virtual System associated with the session |
| Source Zone (from) | Zone the session was sourced from |
| Destination Zone (to) | Zone the session was destined to |
| Ingress Interface (inbound_if) | Interface that the session was sourced from |
| Egress Interface (outbound_if) | Interface that the session was destined to |
| Log Forwarding Profile (logset) | Log Forwarding Profile that was applied to the session |
| Session ID (sessionid) | An internal numerical identifier applied to each session |
| Repeat Count (repeatcnt) | Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only |
| Source Port (sport) | Source port utilized by the session |
| Destination Port (dport) | Destination port utilized by the session |
| NAT Source Port (natsport) | Post-NAT source port |
| NAT Destination Port (natdport) | Post-NAT destination port |
| Flags (flags) | 32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value:<br>• 0x80000000—session has a packet capture (PCAP)<br>• 0x02000000—IPv6 session<br>• 0x01000000—SSL session was decrypted (SSL Proxy)<br>• 0x00800000—session was denied via URL filtering<br>• 0x00400000—session has a NAT translation performed (NAT)<br>• 0x00200000—user information for the session was captured via the captive portal (Captive Portal)<br>• 0x00080000—X-Forwarded-For value from a proxy is in the source user field<br>• 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction)<br>• 0x00008000—session is a container page access (Container Page)<br>• 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above<br>• 0x00000800—symmetric return was used to forward traffic for this session |
| Protocol (proto) | IP protocol associated with the session |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Action (action) | Action taken for the session; values are alert, allow, deny, drop, drop-all-packets, reset-client, reset-server, reset-both, block-url.<br>• Alert—threat or URL detected but not blocked<br>• Allow—flood detection alert<br>• Deny—flood detection mechanism activated and deny traffic based on configuration<br>• Drop—threat detected and associated session was dropped<br>• Drop-all-packets—threat detected and session remains, but drops all packets<br>• Reset-client—threat detected and a TCP RST is sent to the client<br>• Reset-server—threat detected and a TCP RST is sent to the server<br>• Reset-both—threat detected and a TCP RST is sent to both the client and the server<br>• Block-url—URL request was blocked because it matched a URL category that was set to be blocked |
| Miscellaneous (misc) | Field with variable length with a maximum of 1023 characters<br>The actual URI when the subtype is URL<br>File name or file type when the subtype is file<br>File name when the subtype is virus<br>File name when the subtype is WildFire |
| Threat ID (threatid) | Palo Alto Networks identifier for the threat. It is a description string followed by a 64-bit numerical identifier in parentheses for some Subtypes:<br>• 8000 – 8099—scan detection<br>• 8500 – 8599—flood detection<br>• 9999—URL filtering log<br>• 10000 – 19999—sypware phone home detection<br>• 20000 – 29999—spyware download detection<br>• 30000 – 44999—vulnerability exploit detection<br>• 52000 – 52999—filetype detection<br>• 60000 – 69999—data filtering detection<br>• 100000 – 2999999—virus detection<br>• 3000000 – 3999999—WildFire signature feed<br>• 4000000-4999999—DNS Botnet signatures |
| Category (category) | For URL Subtype, it is the URL Category; For WildFire subtype, it is the verdict on the file and is either ‘malicious’ or ‘benign’; For other subtypes, the value is ‘any’. |
| Severity (severity) | Severity associated with the threat; values are informational, low, medium, high, critical |

| Field Name | Description |
|---|---|
| Direction (direction) | Indicates the direction of the attack, *client-to-server* or *server-to-client*<br>• 0—direction of the threat is client to server<br>• 1—direction of the threat is server to client |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama. |
| Source Location (srcloc) | Source country or Internal region for private addresses. Maximum length is 32 bytes. |
| Destination Location (dstloc) | Destination country or Internal region for private addresses. Maximum length is 32 bytes. |
| Content Type (contenttype) | Applicable only when Subtype is URL.<br>Content type of the HTTP response data. Maximum length 32 bytes. |
| PCAP ID (pcap_id) | Pcap-ID is a 64 bit unsigned integral denoting an ID to correlate threat pcap files with extended pcaps taken as a part of that flow. All threat logs will contain either a pcap_id of 0 (no associated pcap), or an ID referencing the extended pcap file. |
| File Digest (filedigest) | Only for WildFire subtype; all other types do not use this field<br>The filedigest string shows the binary hash of the file sent to be analyzed by the WildFire service. |
| Cloud (cloud) | Only for WildFire subtype; all other types do not use this field.<br>The cloud string displays the FQDN of either the WildFire appliance (private) or the WildFire cloud (public) from where the file was uploaded for analysis. |
| New in v6.1! User Agent (user_agent) | Only for the URL Filtering subtype; all other types do not use this field.<br>The User Agent field specifies the web browser that the user used to access the URL, for example Internet Explorer. This information is sent in the HTTP request to the server. |
| New in v6.1! File Type (filetype) | Only for WildFire subtype; all other types do not use this field.<br>Specifies the type of file that the firewall forwarded for WildFire analysis. |
| New in v6.1! X-Forwarded-For (xff) | Only for the URL Filtering subtype; all other types do not use this field.<br>The X-Forwarded-For field in the HTTP header contains the IP address of the user who requested the web page. It allows you to identify the IP address of the user, which is useful particularly if you have a proxy server on your network that replaces the user IP address with its own address in the source IP address field of the packet header. |
| New in v6.1! Referer (referer) | Only for the URL Filtering subtype; all other types do not use this field.<br>The Referer field in the HTTP header contains the URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested. |
| New in v6.1! Sender (sender) | Only for WildFire subtype; all other types do not use this field.<br>Specifies the name of the sender of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall. |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Subject (subject)<br>New in v6.1! | Only for WildFire subtype; all other types do not use this field.<br>Specifies the subject of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall. |
| Recipient (recipient)<br>New in v6.1! | Only for WildFire subtype; all other types do not use this field.<br>Specifies the name of the receiver of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall. |
| Report ID (reportid)<br>New in v6.1! | Only for WildFire subtype; all other types do not use this field.<br>Identifies the analysis request on the WildFire cloud or the WildFire appliance. |

## HIP Match Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source User, Virtual System, Machine name, OS, Source Address, HIP, Repeat Count, HIP Type, FUTURE_USE, FUTURE_USE, Sequence Number, Action Flags

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of HIP match log; unused |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Source User (srcuser) | Username of the user who initiated the session |
| Virtual System (vsys) | Virtual System associated with the HIP match log |
| Machine Name (machinename) | Name of the user's machine |
| OS | The operating system installed on the user's machine or device (or on the client system) |
| Source Address (src) | IP address of the source user |
| HIP (matchname) | Name of the HIP object or profile |
| Repeat Count (repeatcnt) | Number of times the HIP profile matched |
| HIP Type (matchtype) | Whether the hip field represents a HIP object or a HIP profile |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama |

## Config Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Host, Virtual System, Command, Admin, Client, Result, Configuration Path, Sequence Number, Action Flags, Before Change Detail*, After Change Detail*

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of configuration log; unused |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Host (host) | Host name or IP address of the client machine |
| Virtual System (vsys) | Virtual System associated with the configuration log |
| Command (cmd) | Command performed by the Admin; values are add, clone, commit, delete, edit, move, rename, set. |
| Admin (admin) | Username of the Administrator performing the configuration |
| Client (client) | Client used by the Administrator; values are Web and CLI |
| Result (result) | Result of the configuration action; values are Submitted, Succeeded, Failed, and Unauthorized |
| Configuration Path (path) | The path of the configuration command issued; up to 512 bytes in length |
| Sequance Number (seqno) | A 64bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama. |
| Before Change Detail (before_change_detail) New in v6.1! | This field is in custom logs only; it is not in the default format.<br>It contains the full xpath before the configuration change. |
| After Change Detail (after_change_detail) New in v6.1! | This field is in custom logs only; it is not in the default format.<br>It contains the full xpath after the configuration change. |

 © Palo Alto Networks, Inc.

## System Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Virtual System, Event ID, Object, FUTURE_USE, FUTURE_USE, Module, Severity, Description, Sequence Number, Action Flags

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Virtual System (vsys) | Virtual System associated with the configuration log |
| Event ID (eventid) | String showing the name of the event |
| Object (object) | Name of the object associated with the system event |
| Module (module) | This field is valid only when the value of the Subtype field is general. It provides additional information about the sub-system generating the log; values are general, management, auth, ha, upgrade, chassis |
| Severity (severity) | Severity associated with the event; values are informational, low, medium, high, critical |
| Description (opaque) | Detailed description of the event, up to a maximum of 512 bytes |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama |

## Syslog Severity

The syslog severity is set based on the log type and contents.

| Log Type/Severity | Syslog Severity |
|---|---|
| Traffic | Info |
| Config | Info |
| Threat/System—Informational | Info |
| Threat/System—Low | Notice |
| Threat/System—Medium | Warning |

| Log Type/Severity | Syslog Severity |
|---|---|
| Threat/System—High | Error |
| Threat/System—Critical | Critical |

## Custom Log/Event Format

To facilitate the integration with external log parsing systems, the firewall allows you to customize the log format; it also allows you to add custom *Key: Value* attribute pairs. Custom message formats can be configured under **Device > Server Profiles > Syslog > Syslog Server Profile > Custom Log Format**.

To achieve ArcSight Common Event Format (CEF) compliant log formatting, refer to the CEF Configuration Guide.

## Escape Sequences

Any field that contains a comma or a double-quote is enclosed in double quotes. Furthermore, if a double-quote appears inside a field it is escaped by preceding it with another double-quote. To maintain backward compatibility, the Misc field in threat log is always enclosed in double-quotes.

 © Palo Alto Networks, Inc.

# User-ID

User Identification (User-ID) is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses. The following sections describe the Palo Alto Networks User-ID feature and provide instructions on setting up user- and group-based access:


- User-ID Overview
- User-ID Concepts
- Enable User-ID
- Map Users to Groups
- Map IP Addresses to Users
- Configure a Firewall to Share User Mapping Data with Other Firewalls
- Enable User- and Group-Based Policy
- Verify the User-ID Configuration


© Palo Alto Networks, Inc.

# User-ID Overview

User-ID seamlessly integrates Palo Alto Networks firewalls with a range of enterprise directory and terminal services offerings, enabling you to tie application activity and security policies to users and groups—not just IP addresses. In addition, with User-ID enabled, the Application Command Center (ACC), App-Scope, reports, and logs all include usernames in addition to user IP addresses.

The Palo Alto Networks next-generation firewall supports monitoring of the following enterprise services:

- Microsoft Active Directory
- Lightweight Directory Access Protocol (LDAP)
- Novell eDirectory
- Citrix Metaframe Presentation Server or XenApp
- Microsoft Terminal Services

To be able to create policy based on user and group, the firewall must have a list of all available users and their corresponding group mappings that you can select from when defining your policies. It gets this Group Mapping information by connecting directly to your LDAP directory server.

To be able to enforce the user- and group-based policies, the firewall must be able to map the IP addresses in the packets it receives to usernames. User-ID provides many mechanisms to get these IP address-to-username mappings. For example, it uses agents to monitor server logs for login events and/or probe clients, and/or listen for syslog messages from authenticating services. To identify mappings for IP addresses that were not mapped using one of the agent mechanisms, you can configure the firewall to redirect HTTP requests to a captive portal login. You can tailor the mechanisms you use for User Mapping to suit your environment, and you can even use different mechanisms at different sites.

 © Palo Alto Networks, Inc.

**Figure: User-ID**

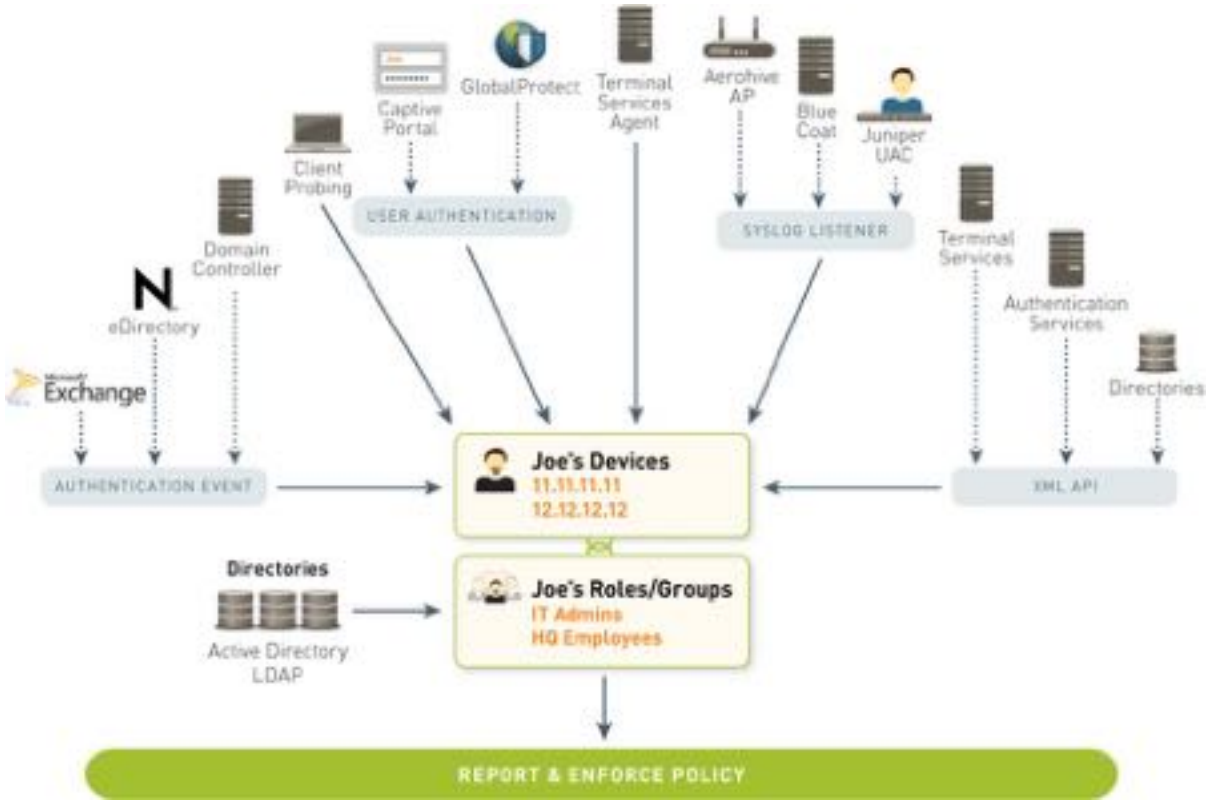


Proceed to User-ID Concepts for information on how User-ID works and Enable User-ID for instructions on setting up User-ID on the firewall.

# User-ID Concepts

▲ Group Mapping
▲ User Mapping

## Group Mapping

In order to define security policies based on user or group, the firewall must retrieve the list of groups and the corresponding list of members from your directory server. To enable this functionality to Map Users to Groups, you must create an LDAP server profile that instructs the firewall how to connect and authenticate to the server and how to search the directory for the user and group information. After you connect to the LDAP server and configure the group mapping functionality for user identification, you will be able to select users or groups when defining your security policies. The firewall supports a variety of LDAP directory servers, including Microsoft Active Directory (AD), Novell eDirectory, and Sun ONE Directory Server.

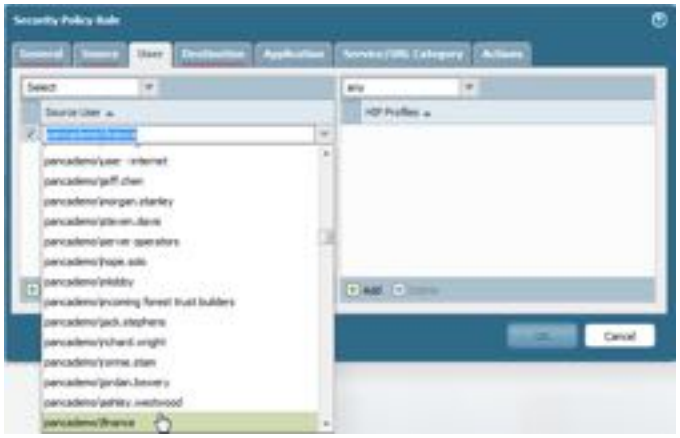

You can then define policies based on group membership rather than on individual users for simplified administration. For example, the following security policy allows access to specific internal applications based on group membership:

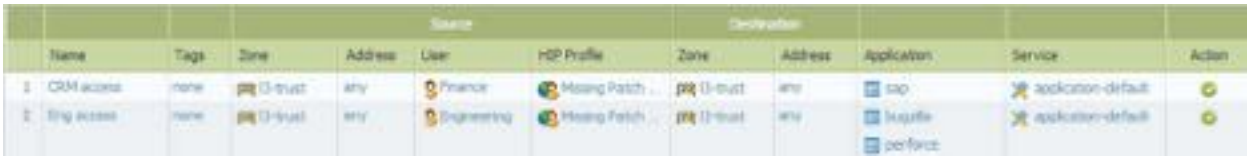

## User Mapping

Having the names of the users and groups is only one piece of the puzzle. The firewall also needs to know which IP addresses map to which users so that security policies can be enforced appropriately. Figure: User-ID illustrates the different methods that are used to identify users and groups on your network and shows how user mapping and group mapping work together to enable user- and group-based security enforcement and visibility.

The following topics describe the different methods of user mapping:

 © Palo Alto Networks, Inc.

- Server Monitoring
- Client Probing
- Port Mapping
- Syslog
- Captive Portal
- GlobalProtect
- User-ID XML API

## Server Monitoring

With server monitoring a User-ID agent—either a Windows-based agent running on a domain server in your network, or the integrated PAN-OS User-ID agent running on the firewall—monitors the security event logs for specified Microsoft Exchange Servers, domain controllers, or Novell eDirectory servers for login events. For example, in an AD environment, you can configure the User-ID agent to monitor the security logs for Kerberos ticket grants or renewals, Exchange server access (if configured), and file and print service connections. Keep in mind that in order for these events to be recorded in the security log, the AD domain must be configured to log successful account login events. In addition, because users can log in to any of the servers in the domain, you must set up server monitoring for all servers in order to capture all user login events.

Because server monitoring requires very little overhead and because the majority of users can generally be mapped using this method, it is recommended as the base user mapping method for most User-ID deployments. See Configure User Mapping Using the Windows User-ID Agent or Configure User Mapping Using the PAN-OS Integrated User-ID Agent for details.

## Client Probing

In a Microsoft Windows environment, you can configure the User-ID agent to probe client systems using Windows Management Instrumentation (WMI). The Windows-based User-ID agent can also perform NetBIOS probing (not supported on the PAN-OS integrated User-ID agent). Probing is particularly useful in environments with a high IP address turnover because changes will be reflected on the firewall more quickly, enabling more accurate enforcement of user-based policies. However, if the correlation between IP addresses and users is fairly static, you probably do not need to enable client probing. Because probing can generate a large amount of network traffic (based on the total number of mapped IP addresses), the agent that will be initiating the probes should be located as close as possible to the end clients.

If probing is enabled, the agent will probe each learned IP address periodically (every 20 minutes by default, but this is configurable) to verify that the same user is still logged in. In addition, when the firewall encounters an IP address for which it has no user mapping, it will send the address to the agent for an immediate probe.

See Configure User Mapping Using the Windows User-ID Agent or Configure User Mapping Using the PAN-OS Integrated User-ID Agent for details.

## Port Mapping

In environments with multi-user systems—such as Microsoft Terminal Server or Citrix environments—many users share the same IP address. In this case, the user-to-IP address mapping process requires knowledge of the source port of each client. To perform this type of mapping, you must install the Palo Alto Networks Terminal Services Agent on the Windows/Citrix terminal server itself to intermediate the assignment of source ports to the various user processes. For terminal servers that do not support the Terminal Services Agent, such as Linux terminal servers, you can use the XML API to send user mapping information from login and logout events to User-ID. See Configure User Mapping for Terminal Server Users for configuration details.

## Syslog

In environments with existing network services that authenticate users—such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms—the firewall User-ID agent (either the Windows agent or the PAN-OS integrated agent on the firewall) can listen for authentication syslog messages from those services. Syslog filters, which are provided by a content update (integrated User-ID agent only) or configured manually, allow the User-ID agent to parse and extract usernames and IP addresses from authentication syslog events generated by the external service, and add the information to the User-ID IP address-to-username mappings maintained by the firewall. See Configure User-ID to Receive User Mappings from a Syslog Sender for configuration details.

**Figure: User-ID Integration with Syslog**

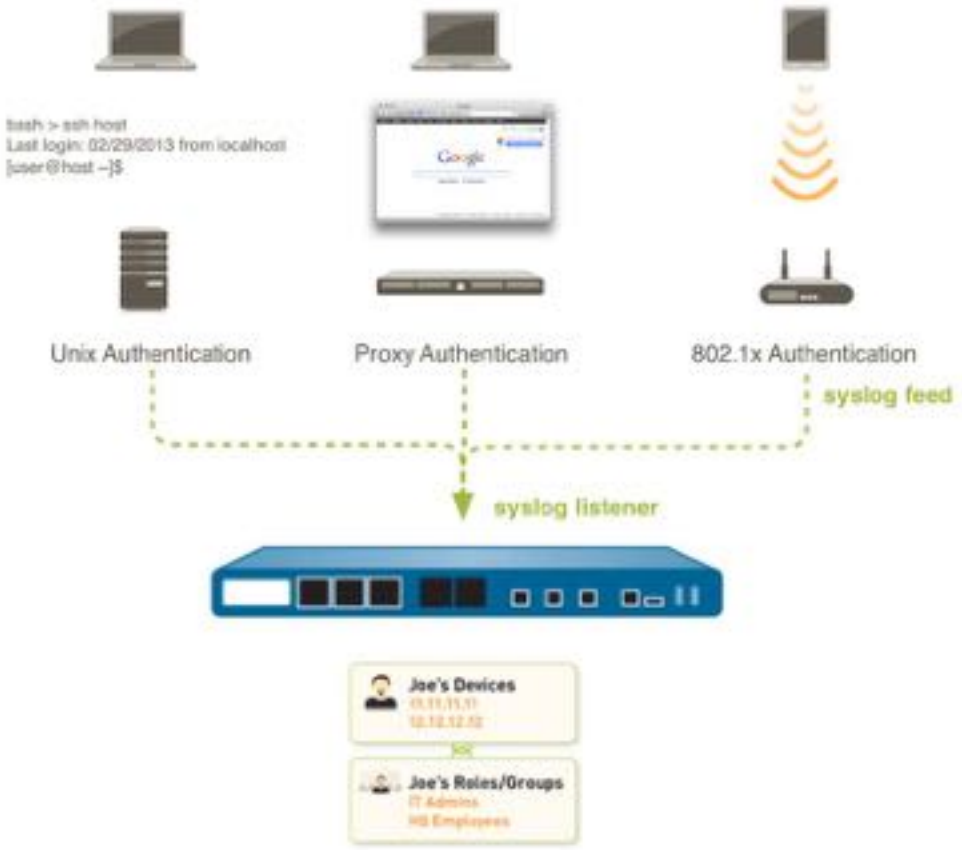


 © Palo Alto Networks, Inc.

## Captive Portal

If the firewall or the User-ID agent are unable to map an IP address to a user—for example if the user is not logged in or is using an operating system such as Linux that is not supported by your domain servers—you can configure Captive Portal. When configured, any web traffic (HTTP or HTTS) matching your Captive Portal policy requires user authentication, either transparently via an NT LAN Manager (NTLM) challenge to the browser, or actively by redirecting the user to a web authentication form for authentication against a RADIUS, LDAP, Kerberos, or local authentication database or using client certificate authentication. See Map IP Addresses to User Names Using Captive Portal for details.

## GlobalProtect

For mobile or roaming users, the GlobalProtect client provides the user mapping information to the firewall directly. In this case, every GlobalProtect user has an agent or app running on the client that requires the user to enter login credentials for VPN access to the firewall. This login information is then added to the User-ID user mapping table on the firewall for visibility and user-based security policy enforcement. Because GlobalProtect users must authenticate to gain access to the network, the IP address-to-username mapping is explicitly known. This is the best solution in sensitive environments where you must be certain of who a user is in order to allow access to an application or service. For more information on setting up GlobalProtect, refer to the GlobalProtect Administrator's Guide.

## User-ID XML API

For other types of user access that cannot be mapped using any of the standard user mapping methods or Captive Portal—for example, to add mappings of users connecting from a third-party VPN solution or users connecting to a 802.1x enabled wireless network—you can use the User-ID XML API to capture login events and send them to the User-ID agent or directly to the firewall. See Send User Mappings to User-ID Using the XML API for details.

© Palo Alto Networks, Inc.

# Enable User-ID

You must complete the following tasks to set up the firewall to user users and groups in policy enforcement, logging, and reporting:

- Map Users to Groups
- Map IP Addresses to Users
- Enable User- and Group-Based Policy
- Verify the User-ID Configuration


© Palo Alto Networks, Inc.

# Map Users to Groups

Use the following procedure to connect to your LDAP directory to enable the firewall to retrieve user-to-group mapping information:

Best practices for group mapping in an Active Directory environment:

- If you have a single domain, you only need one LDAP server profile that connects the firewall to the domain controller with the best connectivity. You can add additional domain controllers for fault tolerance.
- If you have multiple domains and/or multiple forests, you must create a server profile to connect to a domain server in each domain/forest. Take steps to ensure unique usernames in separate forests.
- If you have Universal Groups, create a server profile to connect to the Global Catalogue server.

## Map Users to Groups

Step 1 Create an LDAP Server Profile that specifies how to connect to the directory servers you want the firewall to use to obtain group mapping information.

1. Select **Device > Server Profiles > LDAP**.
2. Click **Add** and then enter a **Name** for the profile.
3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.
4. Click **Add** to add a new LDAP server entry and then enter a **Server** name to identify the server (1-31 characters) and the IP **Address** and **Port** number the firewall should use to connect to the LDAP server (default=389 for LDAP; 636 for LDAP over SSL). You can add up to four LDAP servers to the profile, however, all the servers you add to a profile must be of the same type. For redundancy you should add at least two servers.

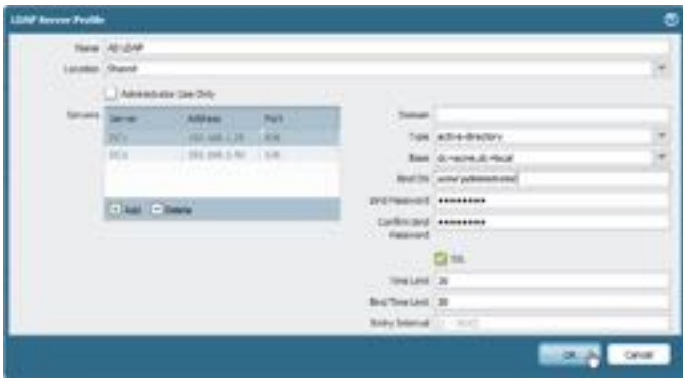

5. Enter the LDAP **Domain** name to prepend to all objects learned from the server. The value you enter here depends on your deployment:
   - If you are using Active Directory, you must enter the NetBIOS domain name; NOT a FQDN (for example, enter `acme`, not `acme.com`). Note that if you need to collect data from multiple domains, you will need to create separate server profiles.
   - If you are using a global catalog server, leave this field blank.
6. Select the **Type** of LDAP server you are connecting to. The correct LDAP attributes in the group mapping settings will automatically be populated based on your selection. However, if you have customized your LDAP schema, you may need to modify the default settings.
7. In the **Base** field, select the DN that corresponds to the point in the LDAP tree where you want the firewall to begin its search for user and group information.
8. Enter the authentication credentials for binding to the LDAP tree in the **Bind DN**, **Bind Password**, and **Confirm Bind Password** fields. The Bind DN can be in either User Principal Name (UPN) format (for example, `administrator@acme.local`) or it can be a fully qualified LDAP name (for example, `cn=administrator,cn=users,dc=acme,dc=local`).
9. If you want the firewall to communicate with the LDAP server(s) over a secure connection, select the SSL check box. If you enable SSL, make sure that you have also specified the appropriate port number.
10. Click **OK** to save the profile.

**Map Users to Groups (Continued)**

Step 2 Add the LDAP server profile to the User-ID Group Mapping configuration.

1. Select **Device > User Identification > Group Mapping Settings** and click **Add**.
2. Enter a unique **Name** to identify the group mapping configuration.
3. Select the **Server Profile** you created in Step 1.
4. (Optional) To filter the groups that User-ID tracks for group mapping, in the Group Objects section enter a **Search Filter** (LDAP query), then specify the **Object Class** (group definition), **Group Name**, and **Group Member**.
5. (Optional) To filter the users that User-ID tracks for group mapping, in the User Objects section enter a **Search Filter** (LDAP query), then specify the **Object Class** (user definition) and **User Name**.
6. (Optional) To match User-ID information with email header information identified in the links and attachments of emails forwarded to WildFire, enter the list of email domains in your organization in the Mail Domains section, **Domain List** field. Use commas to separate multiple domains (up to 256 characters). After you click **OK** (Step 9), PAN-OS automatically populates the **Mail Attributes** field based on your LDAP server type (Sun/RFC, Active Directory, or Novell). When a match occurs, the username in the WildFire log email header section will contain a link that opens the **ACC** tab, filtered by the user or group of users.
7. Make sure the **Enabled** check box is selected.
8. (Optional) To limit which groups appear in security policy, select the **Group Include List** tab and browse the LDAP tree to locate the desired groups. For each group you want to include, select it in the **Available Groups** list and click the add ⊞ icon to move it to the **Included Groups** list.
9. Click **OK** to save the settings.

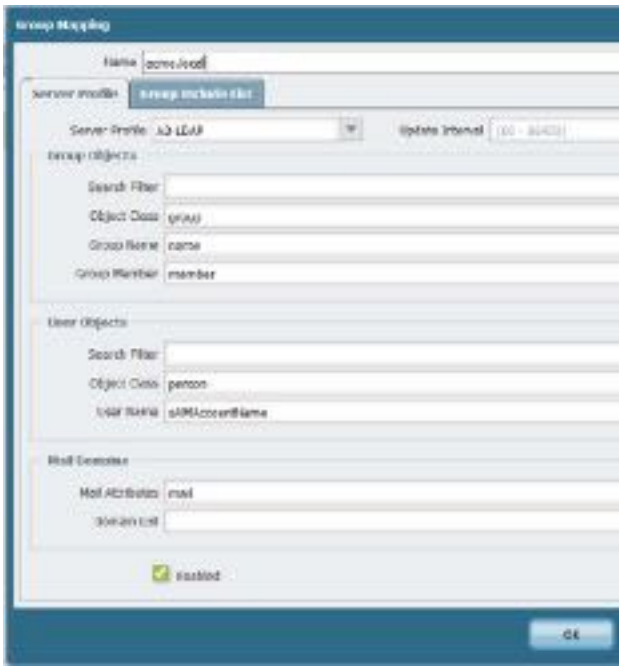


Step 3 **Commit** the configuration.

 © Palo Alto Networks, Inc.

# Map IP Addresses to Users

The tasks you need to perform to map IP addresses to usernames depends on the type and location of the client systems on your network. Complete as many of the following tasks as necessary to enable mapping of your client systems:

- To map users as they log in to your Exchange servers, domain controllers, or eDirectory servers, or Windows clients that can be directly probed, you must configure a User-ID agent to monitor the server logs and/or probe client systems. You can either install the standalone Windows User-ID agent on one or more member servers in the domain that contains the servers and clients to be monitored (see Configure User Mapping Using the Windows User-ID Agent) or you can configure the on-firewall User-ID agent that is integrated with PAN-OS (Configure User Mapping Using the PAN-OS Integrated User-ID Agent). For guidance as to which agent configuration is appropriate for your network and the number and placements of agents that are required, refer to Architecting User Identification Deployments.
- If you have clients running multi-user systems such as Microsoft Terminal Server or Citrix Metaframe Presentation Server or XenApp, see Configure the Palo Alto Networks Terminal Server Agent for User Mapping for instructions on how to install and configure the agent on a Windows server. If you have a multi-user system that is not running on Windows, you can use the User-ID XML API to send IP address to username mappings directly to the firewall. See Retrieve User Mappings from a Terminal Server Using the User-ID XML API.
- To obtain user mappings from existing network services that authenticate users, such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms, configure the User-ID agent (either the Windows agent or the agentless user mapping feature on the firewall) to listen for authentication syslog messages from those services. See Configure User-ID to Receive User Mappings from a Syslog Sender.
- If you have users with client systems that are not logged into your domain servers—for example, users running Linux clients that do not log in to the domain—see Map IP Addresses to User Names Using Captive Portal.
- For other clients that you are unable to map using the previous methods, you can use the User-ID XML API to add user mappings directly to the firewall. See Send User Mappings to User-ID Using the XML API.
- Because policy is local to each firewall, each firewall must have a current list of IP address-to-username mappings in order to accurately enforce security policy by group or user. However, you can configure one firewall to collect all the user mappings and distribute them to the other firewalls. For details, see Configure a Firewall to Share User Mapping Data with Other Firewalls.

© Palo Alto Networks, Inc.

# Configure User Mapping Using the Windows User-ID Agent

In most cases, the majority of your network users will have logins to your monitored domain services. For these users, the Palo Alto Networks User-ID agent monitors the servers for login events and performs the IP address-to-user mapping. The way you configure the User-ID agent depends on the size of your environment and the location of your domain servers. As a best practice, you should locate your User-ID agents near your monitored servers (that is, the monitored servers and the Windows User-ID agent should not be across a WAN link from each other). This is because most of the traffic for user mapping occurs between the agent and the monitored server, with only a small amount of traffic—the delta of IP address mappings since the last update—from the agent to the firewall.

The following topics describe how to install and configure the User-ID Agent and how to configure the firewall to retrieve user mapping information from the agent:

- Install the User-ID Agent
- Configure the User-ID Agent for User Mapping

## Install the User-ID Agent

The following procedure shows how to install the User-ID agent on a member server in the domain and set up the service account with the required permissions. If you are upgrading, the installer will automatically remove the older version, however, it is a good idea to back up the config.xml file before running the installer.

For information about the system requirements for installing the Windows-based User-ID agent and for information on the supported server OS versions are supported, refer to "Operating System (OS) Compatibility User-ID Agent" in the User-ID Agent Release Notes, which are available on the Palo Alto Networks Software Updates page.

| Install the Windows User-ID Agent | | |
|---|---|---|
| Step 1 | Decide where to install the User-ID agent(s).<br><br>The User-ID agent queries the Domain Controller and Exchange server logs using Microsoft Remote Procedure Calls (MSRPCs), which require a complete transfer of the entire log at each query. Therefore, you should always install one or more User-ID agents at each site that has servers to be monitored.<br><br> For more detailed information on where to install User-ID agents, refer to Architecting User Identification (User-ID) Deployments. | • You must install the User-ID agent on a system running one of the supported OS versions: see "Operating System (OS) Compatibility User-ID Agent" in the User-ID Agent Release Notes.<br>• Make sure the system you plan to install the User-ID agent on is a member of the domain that the servers it will be monitoring belong to.<br>• As a best practice, install the User-ID agent close to the servers it will be monitoring (there is more traffic between the User-ID agent and the monitored servers than there is between the User-ID agent and the firewall, so locating the agent close to the monitored servers optimizes bandwidth usage).<br>• To ensure the most comprehensive mapping of users, you must monitor all servers that contain user login information. You may need to install multiple User-ID agents to efficiently monitor all of your resources. |

 © Palo Alto Networks, Inc.

| Install the Windows User-ID Agent (Continued) | |
|---|---|
| Step 2 Download the User-ID agent installer.<br><br>As a best practice, install the User-ID agent version that is the same as the PAN-OS version running on the firewalls. | 1. Log in to Palo Alto Networks Support site.<br>2. Select **Software Updates** from the Manage Devices section.<br>3. Scroll to the User Identification Agent section of the screen and **Download** the version of the User-ID agent you want to install.<br>4. Save the `UaInstall-x.x.x-xx.msi` file on the system(s) where you plan to install the agent. |
| Step 3 Run the installer as an administrator.<br> | 1. To launch a command prompt as an administrator, click Start and right-click **Command Prompt** and then select **Run as administrator**.<br>2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop you would enter the following:<br>`C:\Users\administrator.acme>cd Desktop`<br>`C:\Users\administrator.acme\Desktop>UaInstall-6.0.0-1.msi`<br>3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to the `C:\Program Files (x86)\Palo Alto Networks\User-ID Agent` folder, but you can **Browse** to a different location.<br>4. When the installation completes, **Close** the setup window. |
| Step 4 Launch the User-ID Agent application.<br> 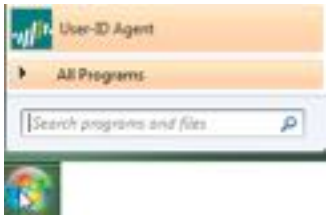 | 1. Click Start and select **User-ID Agent**. |
| Step 5 (Optional) Change the service account that the User-ID agent uses to log in.<br> | By default, the agent uses the administrator account used to install the .msi file. However, you may want to switch this to a restricted account as follows:<br>1. Select **User Identification > Setup** and click **Edit**.<br>2. Select the **Authentication** tab and enter the service account name that you want the User-ID agent to use in the **User name for Active Directory** field.<br>3. Enter the **Password** for the specified account. |

| Install the Windows User-ID Agent (Continued) | |
|---|---|
| Step 6 (Optional) Assign account permissions to the installation folder.<br><br>You only need to perform this step if the service account you configured for the User-ID agent is not a member of the administrators group for the domain or a member of both the Server Operators and the Event Log Readers groups. | 1. Give the service account permissions to the installation folder:<br>a. From the Windows Explorer, navigate to `C:\Program Files\Palo Alto Networks` and right-click the folder and select **Properties**.<br>b. On the **Security** tab, **Add** the User-ID agent service account and assign it permissions to **Modify**, **Read & execute**, **List folder contents**, and **Read** and then click **OK** to save the account settings.<br>2. Give the service account permissions to the User-ID Agent registry sub-tree:<br>a. Run `regedit32` and navigate to the Palo Alto Networks sub-tree in one of the following locations:<br>– **32-bit systems**—`HKEY_LOCAL_MACHINE\Software\Palo Alto Networks`<br>– **64-bit systems**—`HKEY_LOCAL_MACHINE\Software\WOW6432Node\Palo Alto Networks`<br>b. Right-click the Palo Alto Networks node and select **Permissions**.<br>c. Assign the User-ID service account **Full Control** and then click **OK** to save the setting.<br>3. On the domain controller, add the service account to the builtin groups to enable privileges to read the security log events (Event Log Reader group) and open sessions (Server Operator group):<br>a. Run the MMC and Launch the Active Directory Users and Computers snap-in.<br>b. Navigate to the Builtin folder for the domain and then right-click each group you need to edit (Event Log Reader and Server Operator) and select **Add to Group** to open the properties dialog.<br>c. Click **Add** and enter the name of the service account that you configured the User-ID service to use and then click **Check Names** to validate that you have the proper object name.<br>d. Click **OK** twice to save the settings. |

## Configure the User-ID Agent for User Mapping

The Palo Alto Networks User-ID agent is a Windows service that connects to servers on your network—for example, Active Directory servers, Microsoft Exchange servers, and Novell eDirectory servers—and monitors the logs for login events. The agent uses this information to map IP addresses to usernames. Palo Alto Networks firewalls connect to the User-ID agent to retrieve this user mapping information, enabling visibility into user activity by username rather than IP address and enables user- and group-based security enforcement.

 © Palo Alto Networks, Inc.

For information about the server OS versions supported by the User-ID agent, refer to "Operating System (OS) Compatibility User-ID Agent" in the *User-ID Agent Release Notes*, which are available on the Palo Alto Networks Software Updates page.

| **Map IP Addresses to Users Using the User-ID Agent** | | |
|---|---|---|
| Step 1 | Launch the User-ID Agent application. | 1. Select **User-ID Agent** from the Windows Start menu.<br> |
| Step 2 | Define the servers the User-ID agent should monitor to collect IP address to user mapping information.<br>The User-ID agent can monitor a total of up to 100 servers, including syslog senders.<br>Keep in mind that in order to collect all of the required mappings, you must connect to all servers that your users log in to in order to monitor the security log files on all servers that contain login events. | 1. Select **User Identification > Discovery**.<br>2. In the **Servers** section of the screen, click **Add**.<br>3. Enter a **Name** and **Server Address** for the server to be monitored. The network address can be a FQDN or an IP address.<br>4. Select the **Server Type** (**Microsoft Active Directory**, **Microsoft Exchange**, **Novell eDirectory**, or **Syslog Sender**) and then click **OK** to save the server entry. Repeat this step for each server to be monitored.<br>5. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Auto Discover**.<br>The auto-discovery locates domain controllers in the local domain only; you must manually add Exchange servers, eDirectory servers, and syslog senders.<br>e. (Optional) To tune the frequency at which the firewall polls configured servers for mapping information, select **User Identification > Setup** and **Edit** the Setup section. On the **Server Monitor** tab, modify the value in the **Server Log Monitor Frequency (seconds)** field. As a best practice, you should increase the value in this field to 5 seconds in environments with older Domain Controllers or high-latency links. Click **OK** to save the changes. |

| Map IP Addresses to Users Using the User-ID Agent (Continued) | | |
|---|---|---|
| Step 3 | (Optional) If you configured the agent to connect to a Novell eDirectory server, you must specify how the agent should search the directory. | 1. Select **User Identification > Setup** and click **Edit** in the Setup section of the window.<br>2. Select the **eDirectory** tab and then complete the following fields:<br>• **Search Base**—The starting point or root context for agent queries, for example: `dc=domain1, dc=example, dc=com`.<br>• **Bind Distinguished Name**—The account to use to bind to the directory, for example: `cn=admin, ou=IT, dc=domain1, dc=example, dc=com`.<br>• **Bind Password**—The bind account password. The agent saves the encrypted password in the configuration file.<br>• **Search Filter**—The search query for user entries (default is `objectClass=Person`).<br>• **Server Domain Prefix**—A prefix to uniquely identify the user. This is only required if there are overlapping name spaces, such as different users with the same name from two different directories.<br>• **Use SSL**—Select the check box to use SSL for eDirectory binding.<br>• **Verify Server Certificate**—Select the check box to verify the eDirectory server certificate when using SSL. |
| Step 4 | (Optional) Enable client probing.<br>Client probing is useful in environments where IP addresses are not tightly bound to users because it ensures that previously mapped addresses are still valid. However, as the total number of learned IP addresses grows, so does the amount of traffic generated. As a best practice, only enable probing on network segments where IP address turnover is high.<br>For more details on the placement of User-ID agents using client probing, refer to Architecting User Identification (User-ID) Deployments. | 1. On the **Client Probing** tab, select the **Enable WMI Probing** check box and/or the **Enable NetBIOS Probing** check box.<br>2. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client.<br> For NetBIOS probing to work effectively, each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled. WMI probing is always preferred over NetBIOS whenever possible. |
| Step 5 | Save the configuration. | Click **OK** to save the User-ID agent setup settings and then click **Commit** to restart the User-ID agent and load the new settings. |

 © Palo Alto Networks, Inc.

| Map IP Addresses to Users Using the User-ID Agent (Continued) | | |
|---|---|---|
| Step 6 | (Optional) Define the set of users for which you do not need to provide IP address-to-user name mappings, such as kiosk accounts.<br> You can also use the `ignore-user` list to identify users whom you want to force to authenticate using Captive Portal. | Create an `ignore_user_list.txt` file and save it to the User-ID Agent folder on the domain server where the agent is installed.<br>List the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Each user account name must be on a separate line. For example:<br>`SPAdmin`<br>`SPInstall`<br>`TFSReport` |
| Step 7 | Configure the firewalls to connect to the User-ID agent. | Complete the following steps on each firewall you want to connect to the User-ID agent to receive user mappings:<br>1. Select **Device > User Identification > User-ID Agents** and click **Add**.<br>2. Enter a **Name** for the User-ID agent.<br>3. Enter the IP address of the Windows **Host** on which the User-ID Agent is installed.<br>4. Enter the **Port** number on which the agent will be listening for user mapping requests. This value must match the value configured on the User-ID agent. By default, the port is set to 5007 on the firewall and on newer versions of the User-ID agent. However, some older User-ID agent versions use port 2010 as the default.<br>5. Make sure that the configuration is **Enabled** and then click **OK**.<br>6. **Commit** the changes.<br>7. Verify that the **Connected status** displays as connected (a green light). |
| Step 8 | Verify that the User-ID agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent. | 1. Launch the User-ID agent and select **User Identification**.<br>2. Verify that the agent status shows **Agent is running**. If the Agent is not running, click **Start**.<br>3. To verify that the User-ID agent can connect to monitored servers, make sure the Status for each Server is **Connected**.<br>4. To verify that the firewalls can connect to the User-ID agent, make sure the Status for each of the Connected Devices is **Connected**.<br>5. To verify that the User-ID agent is mapping IP addresses to usernames, select **Monitoring** and make sure that the mapping table is populated. You can also **Search** for specific users, or **Delete** user mappings from the list. |

© Palo Alto Networks, Inc.

## Configure User Mapping Using the PAN-OS Integrated User-ID Agent

The following procedure shows how to configure the PAN-OS integrated agent on the firewall for user mapping. The integrated User-ID agent performs the same tasks as the Windows-based agent with the exception of NetBIOS client probing (WMI probing is supported).

| Map IP Addresses to Users Using the Integrated User-ID Agent | | |
|---|---|---|
| Step 1 | Create an Active Directory (AD) account for the firewall agent that has the privilege levels required to log in to each service or host you plan to monitor to collect user mapping data. | • **Windows 2008 or later domain servers**—Add the account to the Event Log Readers group. If you are using the on-device User-ID agent, the account must also be a member of the Distributed COM Users Group.<br>• **Windows 2003 domain servers**—Assign Manage Auditing and Security Logs permissions through group policy.<br>• **WMI probing**—Make sure the account has rights to read the CIMV2 namespace; by default, Domain Administrator and Server Operator accounts have this permission.<br>• **NTLM authentication**—Because the firewall must join the domain if you are using Captive Portal NTLM authentication with an on-device User-ID agent, the Windows account you create for NTLM access must have administrative privileges. Note that due to AD restrictions on virtual systems running on the same host, if you have configured multiple virtual systems, only vsys1 will be able to join the domain. |

 © Palo Alto Networks, Inc.

**Map IP Addresses to Users Using the Integrated User-ID Agent (Continued)**

| | | |
|---|---|---|
| Step 2 | Define the servers the firewall should monitor to collect IP address-to-user mapping information. You can define entries for a total of up to 100 servers of all types, including syslog senders.<br><br>Keep in mind that in order to collect all of the required mappings, you must connect to all servers that your users log in to so that the firewall can monitor the security log files on all servers that contain login events. | 1. Select **Device > User Identification > User Mapping**.<br>2. In the Server Monitoring section of the screen, click **Add**.<br>3. Enter a **Name** and **Network Address** for the server. The network address can be a FQDN or an IP address.<br>4. Select the **Type** of server.<br>5. Make sure the **Enabled** check box is selected and then click **OK**<br>6. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Discover**.<br> <br>The auto-discovery feature is for domain controllers only; you must manually add any Exchange servers or eDirectory servers you want to monitor.<br>7. (Optional) Specify the frequency at which the firewall polls Windows servers for mapping information. This is the interval between the end of the last query and the start of the next query.<br>If the query load is high, the observed delay between queries might significantly exceed the specified frequency.<br>a. In the **Palo Alto Networks User ID Agent Setup** section, click the Edit icon.<br>b. Select the **Server Monitor** tab and specify the **Server Log Monitor Frequency** in seconds (default is 2, range is 1-3600).<br>As a best practice, increase the value in this field to 5 seconds in environments with older domain controllers or high-latency links.<br>c. Click **OK** to save the changes. |
| Step 3 | Set the domain credentials for the account the firewall will use to access Windows resources. This is required for monitoring Exchange servers and domain controllers as well as for WMI probing. | 1. Edit the Palo Alto Networks User ID Agent Setup section of the screen.<br>2. On the **WMI Authentication** tab, enter the **User Name** and **Password** for the account that will be used to probe the clients and monitor servers. Enter the user name using the domain\username syntax.<br> |

| Map IP Addresses to Users Using the Integrated User-ID Agent (Continued) | | |
|---|---|---|
| Step 4 | (Optional) Enable WMI probing.<br><br>The on-device agent does not support NetBIOS probing; it is supported on the Windows-based User-ID agent only. | 1. On the **Client Probing** tab, select the **Enable Probing** check box.<br>2. (Optional) If necessary, modify the **Probe Interval** (in minutes) to ensure it is long enough for the User-ID agent to probe all the learned IP addresses (default is 20, range is 1-1440). This is the interval between the end of the last probe request and the start of the next request.<br><br>If the request load is high, the observed delay between requests might significantly exceed the specified interval.<br>3. Make sure the Windows firewall will allow client probing by adding a a remote administration exception to the Windows firewall for each probed client. |
| Step 5 | Save the configuration. | 1. Click **OK** to save the User-ID agent setup settings.<br>2. Click **Commit** to save the configuration. |
| Step 6 | (Optional) Define the set of users for which you do not need to provide IP address-to-user name mappings, such as kiosk accounts.<br><br>You can also use the `ignore-user` list to identify users whom you want to force to authenticate using Captive Portal. | 1. Open a CLI session to the firewall.<br>2. To add the list of user accounts for which you do not want the firewall to perform mapping, run the following command:<br>**set user-id-collector ignore-user <value>**<br>where <value> is a list of the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Separate entries with a space and do not include the domain name with the username. For example:<br>`set user-id-collector ignore-user SPAdmin SPInstall TFSReport`<br>3. **Commit** your changes. |
| Step 7 | Verify the configuration. | 1. From the CLI, enter the following command:<br>`show user server-monitor state all`<br>2. On the **Device > User Identification > User Mapping** tab in the web interface, verify that the **Status** of each server you configured for server monitoring is **Connected**. |

 © Palo Alto Networks, Inc.

# Configure User-ID to Receive User Mappings from a Syslog Sender

The following topics describe how to configure the User-ID agent (either the Windows agent or the integrated agent on the firewall) as a Syslog listener:

- Configure the Integrated User-ID Agent as a Syslog Listener
- Configure the Windows User-ID Agent as a Syslog Listener

## Configure the Integrated User-ID Agent as a Syslog Listener

The following workflow describes how to configure the PAN-OS integrated User-ID agent to receive syslog messages from authenticating services.

The PAN-OS integrated User-ID agent accepts syslogs over SSL and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages. However, if you must use UDP, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending UDP traffic to the firewall.

| **Collect User Mappings from Syslog Senders** | |
|---|---|
| Step 1 Determine whether there is a pre-defined syslog filter for your particular syslog sender(s).<br><br>Palo Alto Networks provides several pre-defined syslog filters, which are delivered as Application content updates and are therefore updated dynamically as new filters are developed. The pre-defined filters are global to the firewall, whereas manually-defined filters apply to a single virtual system only.<br><br> Any new syslog filters in a given content release will be documented in the corresponding release note along with the specific regex used to define the filter. | 1. Verify that your Application or Application and Threat database is up to date:<br>a. Select **Device > Dynamic Updates**.<br>b. Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates.<br>c. If a new update is available, **Download** and **Install** it.<br>2. Check to see what pre-defined filters are available:<br>a. Select **Device > User Identification > User Mapping**.<br>b. In the Server Monitoring section of the screen, click **Add**.<br>c. Select **Syslog Sender** as the server **Type**.<br>d. Select the **Filter** drop-down and check to see if there is a filter for the manufacturer and product you plan to forward syslogs from. If the filter you need is available, skip to Step 5 for instructions on defining the servers. If the filter you need is not available, continue to Step 2. |

**Collect User Mappings from Syslog Senders (Continued)**

Step 2 Manually define syslog filter(s) for extracting the User-ID IP address to username mapping information from syslog messages.

In order to be parsed by the User-ID agent, syslog messages must meet the following criteria:

- Each syslog message must be a single-line text string. Line breaks are delimited by a carriage return and a new line (`\r\n`) or a new line (`\n`).
- The maximum allowed size of an individual syslog message is 2048 bytes.
- Syslog messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets.
- A single packet may contain multiple syslog messages.

1. Review the syslogs generated by the authenticating service to identify the syntax of the login events. This enables you to create the matching patterns that will allow the firewall to identify and extract the authentication events from the syslogs.

   While reviewing the syslogs, also determine whether the domain name is included in the log entries. If the authentication logs do not contain domain information, consider defining a default domain name when adding the syslog sender to the monitored servers list in Step 5.
2. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User-ID Agent Setup section.
3. On the **Syslog Filters** tab, **Add** a new syslog parse profile.
4. Enter a name for the **Syslog Parse Profile**.
5. Specify the **Type** of parsing to use to filter out the user mapping information by selecting one of the following options:
   - **Regex Identifier**—With this type of parsing, you specify regular expressions to describe search patterns for identifying and extracting user mapping information from syslog messages. Continue to Step 3 for instructions on creating the regex identifiers.
   - **Field Identifier**—With this type of parsing, you specify a string to match the authentication event, and prefix and suffix strings to identify the user mapping information in the syslogs. Continue to Step 4 for instructions on creating the field identifiers.

 © Palo Alto Networks, Inc.

**Collect User Mappings from Syslog Senders (Continued)**

Step 3 If you selected **Regex Identifier** as the parsing **Type**, create the regex matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a regex configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter, you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event Regex** field. For example, when matched against the example syslog message, the following regex instructs the firewall to extract the first `{1}` instance of the string `authentication success`. The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character: `(authentication\ success){1}`.
2. Enter the regex for identifying the beginning of the username in the authentication success messages in the **Username Regex** field. For example, the regex `User:([a-zA-Z0-9\\\._]+)` would match the string `User:johndoe1` in the example message and extract `acme\johndoe1` as the User-ID.

   If the syslogs do not contain domain information and you require domain names in your user mappings, be sure to enter the **Default Domain Name** when defining the monitored server entry in Step 5.
3. Enter the regex for identifying the IP address portion of the authentication success messages in the **Address Regex** field. For example, the following regular expression `Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})` would match an IPv4 address (`Source:192.168.0.212` in the example syslog).
4. Click **OK**.

**Collect User Mappings from Syslog Senders (Continued)**

Step 4 If you selected **Field Identifier** as the parsing **Type**, define the string matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a field identifier configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter, you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event String** field. For example, when matched against the sample syslog message, you would enter the string `authentication success` to identify authentication events in the syslog.
2. Enter the matching string for identifying the beginning of the username field within the authentication syslog message in the **Username Prefix** field. For example, the string `User:` identifies the beginning of the username field in the sample syslog.
3. Enter the **Username Delimiter** to mark the end of the username field within an authentication syslog message. For example, if the username is followed by a space, you would enter `\s` to indicate that the username field is delimited by a standalone space in the sample log.
4. Enter the matching string for identifying the beginning of the IP address field within the authentication event log in the **Address Prefix** field. For example, the string `Source:` identifies the beginning of the address field in the example log.
5. Enter the **Address Delimiter** to mark the end of the IP address field within the authentication success message within the field. For example, if the address is followed by a line break, you would enter `\n` to indicate that the address field is delimited by a new line.
6. Click **OK**.

Step 5 Define the servers that will send syslog messages to the firewall for user mapping purposes.

You can define a total of up to 100 monitored servers of all types, including syslog senders. The firewall will discard any syslog messages received from servers that are not on this list.

A Syslog sender using SSL to connect will only show a **Status** of **Connected** when there is an active SSL connection. Syslog senders using UDP will not show a **Status** value.

1. Select **Device > User Identification > User Mapping**.
2. In the Server Monitoring section of the screen, click **Add**.
3. Enter a **Name** and **Network Address** for the server.
4. Select **Syslog Sender** as the server **Type**.
5. Make sure the **Enabled** check box is selected.
6. (Optional) If the syslogs that the authenticating device sends do not include domain information in the login event logs, enter the **Default Domain Name** to append to the user mappings.
7. Click **OK** to save the settings.

 © Palo Alto Networks, Inc.

| Collect User Mappings from Syslog Senders (Continued) | |
|---|---|
| Step 6 Enable syslog listener services in the management profile associated with the interface used for user mapping.<br><br>Use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages when using agentless User Mapping on a firewall. However, if you must use UDP, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending UDP traffic to the firewall. | 1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.<br>2. Select **User-ID Syslog Listener-SSL** and/or **User-ID Syslog Listener-UDP**, depending on the protocols you defined when you set up your Syslog Senders in the Server Monitor list.<br><br>On the Windows User-ID agent, the default listening port for syslog over UDP or TCP is 514, but the port value is configurable. For the agentless User Mapping feature on the firewall, only syslog over UDP and SSL are supported and the listening ports (514 for UDP and 6514 for SSL) are not configurable; they are enabled through the management service only.<br><br> <br>3. Click **OK** to save the interface management profile.<br><br>Even after enabling the User-ID Syslog Listener service on the interface, the interface will only accept syslog connections from servers that have a corresponding entry in the User-ID monitored servers configuration. Connections or messages from servers that are not on the list will be discarded. |
| Step 7 Save the configuration. | Click **Commit** to save the configuration. |

**Collect User Mappings from Syslog Senders (Continued)**

Step 8 Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
        UDP Syslog Listener Service is enabled
        SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)     Host: Syslog2(10.5.204.41)
        number of log messages                          : 1000
        number of auth. success messages                : 1000
        number of active connections                    : 0
        total connections made                          : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name                            TYPE     Host            Vsys    Status

------------------------------------------------------------------------------
AD                              AD       10.2.204.43     vsys1   Connected

Syslog Servers:
Name                       Connection Host            Vsys    Status
------------------------------------------------------------------------------
Syslog1                    UDP        10.5.204.40     vsys1   N/A
Syslog2                    SSL        10.5.204.41     vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP              Vsys   From    User                              IdleTimeout(s) M
axTimeout(s)
--------------- ------ ------- --------------------------------- -------------- -
192.168.3.8     vsys1  SYSLOG  acme\jreddick                     2476           2
476
192.168.5.39    vsys1  SYSLOG  acme\jdonaldson                   2480           2
480
192.168.2.147   vsys1  SYSLOG  acme\ccrisp                       2476           2
476
192.168.2.175   vsys1  SYSLOG  acme\jjaso                        2476           2
476
192.168.4.196   vsys1  SYSLOG  acme\jblevins                     2480           2
480
192.168.4.103   vsys1  SYSLOG  acme\bmoss                        2480           2
480
192.168.2.193   vsys1  SYSLOG  acme\esogard                      2476           2
476
192.168.2.119   vsys1  SYSLOG  acme\acallaspo                    2476           2
476
192.168.3.176   vsys1  SYSLOG  acme\jlowrie                      2478           2
478

Total: 9 users
```

## Configure the Windows User-ID Agent as a Syslog Listener

The following workflow describes how to configure a Windows-based User-ID agent to listen for syslogs from authenticating services.

The Windows User-ID agent accepts syslogs over TCP and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, use TCP instead of UDP. In either case, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

 © Palo Alto Networks, Inc.

| Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders | |
|---|---|
| Step 1 Launch the User-ID Agent application. | 1. Click Start and select **User-ID Agent**.<br> |
| Step 2 Manually define syslog filter(s) for extracting the User-ID IP address to username mapping information from syslog messages.<br><br>In order to be parsed by the User-ID agent, syslog messages must meet the following criteria:<br>• Each syslog message must be a single-line text string. Line breaks are delimited by a carriage return and a new line (\r\n) or a new line (\n).<br>• The maximum allowed size of an individual syslog message is 2048 bytes.<br>• Syslog messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets.<br>• A single packet may contain multiple syslog messages. | 1. Review the syslogs generated by the authenticating service to identify the syntax of the login events. This enables you to create the matching patterns that will allow the firewall to identify and extract the authentication events from the syslogs.<br> While reviewing the syslogs, also determine whether the domain name is included in the log entries. If the authentication logs do not contain domain information, consider defining a default domain name when adding the syslog sender to the monitored servers list in Step 5.<br>2. Select **User Identification > Setup** and click **Edit** in the Setup section of the dialog.<br>3. On the **Syslog** tab, **Add** a new syslog parse profile.<br>4. Enter a **Profile Name** and **Description**.<br>5. Specify the **Type** of parsing to use to filter out the user mapping information by selecting one of the following options:<br>• **Regex**—With this type of parsing, you specify regular expressions to describe search patterns for identifying and extracting user mapping information from syslog messages. Continue to Step 3 for instructions on creating the regex identifiers.<br>• **Field**—With this type of parsing, you specify a sting to match the authentication event, and prefix and suffix strings to identify the user mapping information in the syslogs. Continue to Step 4 for instructions on creating the field identifiers. |

**Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)**

Step 3 If you selected **Regex** as the parsing **Type**, create the regex matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a regex configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter, you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event Regex** field. For example, when matched against the example syslog message, the following regex instructs the firewall to extract the first `{1}` instance of the string `authentication success`. The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character: `(authentication\ success){1}`.
2. Enter the regex for identifying the beginning of the username in the authentication success messages in the **Username Regex** field. For example, the regex `User:([a-zA-Z0-9\\\._]+)` would match the string `User:johndoe1` in the example message and extract `acme\johndoe1` as the User-ID.

   If the syslogs do not contain domain information and you require domain names in your user mappings, be sure to enter the **Default Domain Name** when defining the monitored server entry in Step 5.
3. Enter the regex for identifying the IP address portion of the authentication success messages in the **Address Regex** field. For example, the following regular expression `Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})` would match an IPv4 address (`Source:192.168.0.212` in the example syslog).
4. Click **OK** to save the profile.


© Palo Alto Networks, Inc.

**Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)**

**Step 4** If you selected **Field Identifier** as the parsing **Type**, define the string matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a field identifier configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter, you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event String** field. For example, when matched against the sample syslog message, you would enter the string `authentication success` to identify authentication events in the syslog.
2. Enter the matching string for identifying the beginning of the username field within the authentication syslog message in the **Username Prefix** field. For example, the string `User:` identifies the beginning of the username field in the sample syslog.
3. Enter the **Username Delimiter** to mark the end of the username field within an authentication syslog message. For example, if the username is followed by a space, you would enter `\s` to indicate that the username field is delimited by a standalone space in the sample log.
4. Enter the matching string for identifying the beginning of the IP address field within the authentication event log in the **Address Prefix** field. For example, the string `Source:` identifies the beginning of the address field in the example log.
5. Enter the **Address Delimiter** to mark the end of the IP address field within the authentication success message within the field. For example, if the address is followed by a line break, you would enter `\n` to indicate that the address field is delimited by a new line.
6. Click **OK** to save the profile.

**Step 5** Enable the syslog listening service on the agent.

As a best practice, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

1. Select the **Enable Syslog Service** check box.
2. (Optional) Modify the **Syslog Service Port** number to match the port number used by the syslog sender (default=514).
3. To save the agent syslog configuration, click **OK**.

**Step 6** Define the servers that will send syslog messages to the User-ID agent.

You can define up to 100 syslog senders. The User-ID agent will discard any syslog messages received from servers that are not on this list.

1. Select **User Identification > Discovery**.
2. In the **Servers** section of the screen, click **Add**.
3. Enter a **Name** and **Server Address** for the server that will send syslogs to the agent.
4. Select **Syslog Sender** as the **Server Type**.
5. Select a **Filter** you defined in Step 2.
6. (Optional) If the syslogs that the authenticating device sends do not include domain information in the login event logs, enter the **Default Domain Name** to append to the user mappings.
7. Click **OK** to save the settings.

**Step 7** Save the configuration.

Click **Commit** to save the configuration.

**Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)**

Step 8 Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
        UDP Syslog Listener Service is enabled
        SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)     Host: Syslog2(10.5.204.41)
        number of log messages                          : 1000
        number of auth. success messages                : 1000
        number of active connections                    : 0
        total connections made                          : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name                            TYPE     Host            Vsys    Status

------------------------------------------------------------------------------
AD                              AD       10.2.204.43     vsys1   Connected

Syslog Servers:
Name                      Connection Host            Vsys    Status
------------------------------------------------------------------------------
Syslog1                   UDP        10.5.204.40     vsys1   N/A
Syslog2                   SSL        10.5.204.41     vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP              Vsys   From    User                            IdleTimeout(s) M
axTimeout(s)
--------------- ------ ------- ------------------------------- -------------- -
192.168.3.8     vsys1  SYSLOG  acme\jreddick                   2476           2
476
192.168.5.39    vsys1  SYSLOG  acme\jdonaldson                 2480           2
480
192.168.2.147   vsys1  SYSLOG  acme\ccrisp                     2476           2
476
192.168.2.175   vsys1  SYSLOG  acme\jjaso                      2476           2
476
192.168.4.196   vsys1  SYSLOG  acme\jblevins                   2480           2
480
192.168.4.103   vsys1  SYSLOG  acme\bmoss                      2480           2
480
192.168.2.193   vsys1  SYSLOG  acme\esogard                    2476           2
476
192.168.2.119   vsys1  SYSLOG  acme\acallaspo                  2476           2
476
192.168.3.176   vsys1  SYSLOG  acme\jlowrie                    2478           2
478

Total: 9 users
```

 © Palo Alto Networks, Inc.

# Map IP Addresses to User Names Using Captive Portal

If the firewall receives a request from a zone that has User-ID enabled and the source IP address does not have any user data associated with it yet, it checks its Captive Portal policy for a match to determine whether to perform authentication. This is useful in environments where you have clients that are not logged in to your domain servers, such as Linux clients. This user mapping method is triggered only for web traffic (HTTP or HTTPS) that matches a security rule/policy, but that has not been mapped using a different method.

- Captive Portal Authentication Methods
- Captive Portal Modes
- Configure Captive Portal

## Captive Portal Authentication Methods

Captive Portal uses the following methods to obtain user data from the client when a request matches a Captive Portal policy:

| Authentication Method | Description |
|---|---|
| NT LAN Manager (NTLM) | The firewall uses an encrypted challenge-response mechanism to obtain the user's credentials from the browser. When configured properly, the browser will provide the credentials to the firewall transparently without prompting the user, but will display a prompt for credentials if necessary. If the browser cannot perform NTLM or if NTLM authentication fails, the firewall falls back to web form or client certificate authentication, depending on your Captive Portal configuration.<br>By default, IE supports NTLM. Firefox and Chrome can be configured to use it. You cannot use NTLM to authenticate non-Windows clients. |
| Web Form | Requests are redirected to a web form for authentication. You can configure Captive Portal to use a local user database, RADIUS, LDAP, or Kerberos to authenticate users. Although users will always be prompted for credentials, this authentication method works with all browsers and operating systems. |
| Client Certificate Authentication | Prompts the browser to present a valid client certificate for authenticating the user. To use this method, you must provision client certificates on each user system and install the trusted CA certificate used to issue those certificates on the firewall. This is the only authentication method that enables transparent authentication for Mac OS and Linux clients. |

## Captive Portal Modes

The Captive Portal mode defines how web requests are captured for authentication:

© Palo Alto Networks, Inc.

| Mode | Description |
|---|---|
| Transparent | The firewall intercepts the browser traffic per the Captive Portal rule and impersonates the original destination URL, issuing an HTTP 401 to invoke authentication. However, because the firewall does not have the real certificate for the destination URL, the browser will display a certificate error to users attempting to access a secure site. Therefore you should only use this mode when absolutely necessary, such as in Layer 2 or virtual wire deployments. |
| Redirect | The firewall intercepts unknown HTTP or HTTPS sessions and redirects them to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to perform authentication. This is the preferred mode because it provides a better end-user experience (no certificate errors). However, it does require additional Layer 3 configuration. Another benefit of the Redirect mode is that it provides for the use of session cookies, which enable the user to continue browsing to authenticated sites without requiring re-mapping each time the time outs expire. This is especially useful for users who roam from one IP address to another (for example, from the corporate LAN to the wireless network) because they will not need to re-authenticate upon IP address change as long as the session stays open. In addition, if you plan to use NTLM authentication, you must use Redirect mode because the browser will only provide credentials to trusted sites. |

## Configure Captive Portal

The following procedure shows how to configure Captive Portal using the PAN-OS integrated User-ID agent to redirect requests that match a Captive Portal policy to a Layer 3 interface on the firewall.

If you plan to use Captive Portal without using the other User-ID functions (user mapping and group mapping), you do not need to configure an agent.

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent | | |
|---|---|---|
| Step 1 | Make sure the firewall has a route to the servers it will monitor to gather user data (for example, your Domain Controllers and your Exchange servers). | The firewall must be able to communicate with the servers over a designated interface (the management [MGT] interface by default), so make sure that the network where your directory servers reside is accessible from that interface. If this configuration doesn't work in your environment, you must configure Captive Portal using the Window-based User-ID agent. |
| Step 2 | Make sure DNS is configured to resolve your Domain Controller addresses. | To verify proper resolution, ping the server FQDN. For example:<br>`admin@PA-200> ping host dc1.acme.com` |

 © Palo Alto Networks, Inc.

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | | |
|---|---|---|
| Step 3 | (Redirect mode only) Create a Layer 3 interface to which to redirect Captive Portal requests. | 1. Create a management profile to enable the interface to display Captive Portal response pages:<br>a. Select **Network > Interface Mgmt** and click **Add**.<br>b. Enter a **Name** for the profile, select **Response Pages**, and then click **OK**.<br>2. Create the Layer 3 interface. Be sure to attach the management profile you just created (on the **Advanced > Other Info** tab of the Ethernet Interface dialog).<br>3. Create a DNS "A" record that maps the IP address you configured on the Layer 3 interface to an intranet host name (that is, a hostname that does not have a dot in the name, such as `ntlmhost`). |
| Step 4 | (Redirect mode only) To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface to which you are redirecting requests.You can either generate a self-signed certificate or import a certificate that is signed by an external CA.<br> When setting up Captive Portal for the first time, imported certificates may not work. If you plan to use an imported certificate, complete the initial configuration without specifying a **Server Certificate**. After you get Captive Portal working, you can go back and switch to the imported certificate. | To use a self-signed certificate, you must first create a root CA certificate and then use that CA to sign the certificate you will use for Captive Portal as follows:<br>1. To create a root CA certificate, select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**. Enter a **Certificate Name**, such as RootCA. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed). Make sure you select the **Certificate Authority** check box and then click **Generate** the certificate.<br>2. To create the certificate to use for Captive Portal, click **Generate**. Enter a **Certificate Name** and enter the DNS name of the intranet host for the interface as the **Common Name**. In the **Signed By** field, select the CA you created in the previous step. Add an IP address attribute and specify the IP address of the Layer 3 interface to which you will be redirecting requests. **Generate** the certificate.<br> 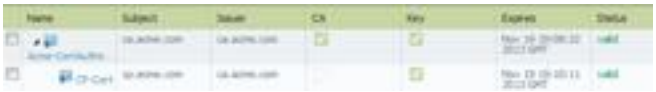 <br>3. To configure clients to trust the certificate, select the CA certificate on the **Device Certificates** tab and click **Export**. You must then import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO). |

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | | |
|---|---|---|
| Step 5 | Set up an authentication mechanism to use when the web form is invoked. Note that even if you plan to use NTLM, you must also set up a secondary authentication mechanism that can be used if NTLM authentication fails or if the user agent does not support it.<br><br>**Best Practices:**<br>•If using RADIUS to authenticate users from the web form, be sure to enter a RADIUS domain. This will be used as the default domain if users don't supply one upon login.<br>•If using AD to authenticate users from the web form, make sure to enter **sAMAccountName** as the **LogonAttribute**. | 1. Configure the firewall to connect to the authentication service you plan to use so that it can access the authentication credentials.<br>• If you plan to authenticate using LDAP, Kerberos, or RADIUS, you must create a server profile that instructs the firewall how to connect to the service and access the authentication credentials for your users. Select **Device > Server Profiles** and add a new profile for the specific service you will be accessing.<br>• If you plan to use local database authentication, you must first create the local database. Select **Device > Local User Database** and add the users and groups to be authenticated.<br>2. Create an authentication profile that references the server profile or local user database you just created. Select **Device > Authentication Profile** and add a new profile for use with Captive Portal. For details on creating a specific type of authentication profile, refer to the online help. |
| Step 6 | (Optional) Set up client certificate authentication. Note that you do not need to set up both an authentication profile and a client certificate profile to enable Captive Portal. If you configure both, the user will be required to authenticate using both methods.<br><br>For details on other certificate profile fields, such as whether to use CRL or OCSP, refer to the online help. | 1. Use a root CA certificate to generate a client certificate for each user who will authenticate to Captive Portal. The CA in this case is usually your enterprise CA, not the firewall.<br>2. Export the CA certificate in Base64 format to a system that the firewall can access.<br>3. Import the root CA certificate from the CA that generated the client certificates onto the firewall:<br>a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.<br>c. **Browse** to the **Certificate File** you downloaded from the CA.<br>d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.<br>e. Select the certificate you just imported on the **Device Certificates** tab to open it.<br>f. Select **Trusted Root CA** and then click **OK**.<br>4. Create the client certificate profile that you will use when you configure Captive Portal.<br>a. Select **Device > Certificates > Certificate Management > Certificate Profile**, click **Add**, and enter a profile **Name**.<br>b. In the **Username Field** drop-down, select the certificate field that contains the user's identity information.<br>c. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you just imported and then click **OK**. |

 © Palo Alto Networks, Inc.

**Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)**

Step 7 Enable NT LAN Manager (NTLM) authentication.

When using the on-device User-ID agent, the firewall must be able to successfully resolve the DNS name of your Domain Controller in order for the firewall to join the domain. The credentials you supply here will be used to join the firewall to the domain upon successful DNS resolution.

1. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User ID Agent Setup section of the screen.
2. On the **NTLM** tab, select the **Enable NTLM authentication processing** check box.
3. Enter the **NTLM Domain** against which the User-ID agent on the firewall should check NTLM credentials.
4. Enter the **Admin User Name** and **Password** for the Active Directory account you created in Step 1 in Map IP Addresses to Users Using the Integrated User-ID Agent for NTLM authentication.

   Do not include the domain in the **Admin User Name** field. Otherwise, the firewall will fail to join the domain.

   Palo Alto Networks recommends that you use a User-ID agent account that is separate from your firewall administrator account.
5. You don't need to configure any other settings for the User-ID agent: click **OK**.

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | |
|---|---|
| Step 8 Configure the Captive Portal settings. | 1. Select **Device > User Identification > Captive Portal Settings** and Edit the **Captive Portal** section of the screen.<br>2. Make sure the **Enabled** check box is selected.<br>3. Set the **Mode**. This example shows how to set up **Redirect** mode.<br>4. (Redirect mode only) Select the **Server Certificate** the firewall should use to redirect requests over SSL. This is the certificate you created in Step 4.<br>5. (Redirect mode only) Specify the **Redirect Host**, which is the intranet hostname that resolves to the IP address of the Layer 3 interface to which you are redirecting requests, as specified in Step 3.<br>6. Select the authentication method to use if NTLM fails (or if you are not using NTLM):<br>• If you are using LDAP, Kerberos, RADIUS, or local database authentication, select the **Authentication Profile** you created in Step 5.<br>• If you are using client certificate authentication, select the **Certificate Profile** you created in Step 6.<br> 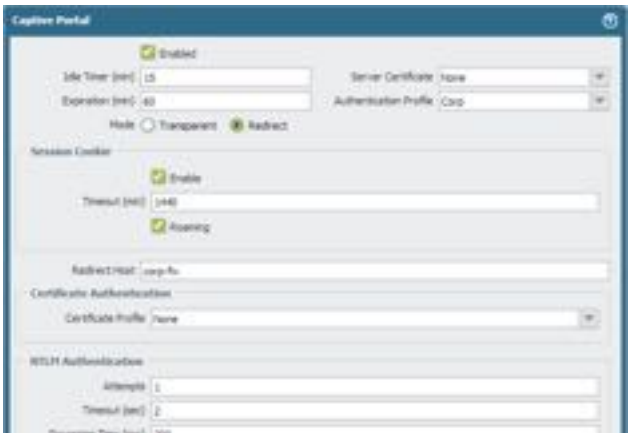 <br>7. Click **OK** to save your settings.<br>8. Click **Commit** to save the Captive Portal configuration. |

 © Palo Alto Networks, Inc.

# Configure User Mapping for Terminal Server Users

Individual terminal server users appear to have the same IP address and therefore an IP address-to-username mapping is not sufficient to identify a specific user. To enable identification of specific users on Windows-based terminal servers, the Palo Alto Networks Terminal Services agent (TS agent) allocates a port range to each user. It then notifies every connected firewall about the allocated port range, which allows the firewall to create an IP address-port-user mapping table and enable user- and group-based security policy enforcement. For non-Windows terminal servers, you can configure the User-ID XML API to extract user mapping information.

The following sections describe how to configure user mapping for terminal server users:

- Configure the Palo Alto Networks Terminal Server Agent for User Mapping
- Retrieve User Mappings from a Terminal Server Using the User-ID XML API

## Configure the Palo Alto Networks Terminal Server Agent for User Mapping

Use the following procedure to install the TS agent on the terminal server. You must install the TS agent on all terminal servers that your users log in to in order to successfully map all your users.

For information about the supported terminal servers supported by the TS Agent, refer to "Operating System (OS) Compatibility TS Agent" in the Terminal Services Agent Release Notes, which are available on the Palo Alto Networks Software Updates page.

| **Install the Windows Terminal Server Agent** | |
|---|---|
| Step 1 Download the TS Agent installer. | 1. Log in to the Palo Alto Networks Support site.<br>2. Select **Software Updates** from the Manage Devices section.<br>3. Scroll to the **Terminal Services Agent** section and **Download** the version of the agent you want to install.<br>4. Save the `TaInstall64.x64-x.x.x-xx.msi` or `TaInstall-x.x.x-xx.msi` file (be sure to select the appropriate version based on whether the Windows system is running a 32-bit OS or a 64-bit OS) on the system(s) where you plan to install the agent. |

**Install the Windows Terminal Server Agent (Continued)**

| | | |
|---|---|---|
| Step 2 | Run the installer as an administrator.<br> | 1. To launch a command prompt as an administrator, click Start and right-click **Command Prompt** and then select **Run as administrator**.<br>2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop you would enter the following:<br>`C:\Users\administrator.acme>cd Desktop`<br>`C:\Users\administrator.acme\Desktop>TaInstall-6.0.0-1.msi`<br>3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to the `C:\Program Files (x86)\Palo Alto Networks\Terminal Server Agent` folder, but you can **Browse** to a different location.<br>4. When the installation completes, **Close** the setup window.<br>If you are upgrading to a TS Agent version that has a newer driver than the existing installation, the installation wizard prompts you to reboot the system after upgrading in order to use the new driver. |
| Step 3 | Launch the Terminal Server Agent application. | Click Start and select **Terminal Server Agent**.<br> 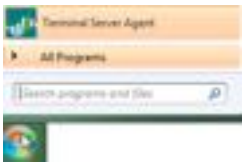 |

 © Palo Alto Networks, Inc.

**Install the Windows Terminal Server Agent (Continued)**

Step 4 Define the range of ports for the TS Agent to allocate to end users.

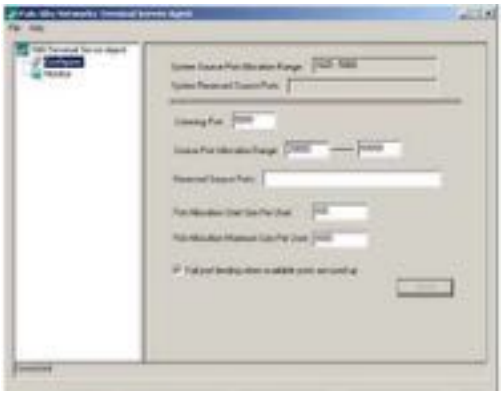

The **System Source Port Allocation Range** and **System Reserved Source Ports** fields specify the range of ports that will be allocated to non-user sessions. Make sure the values specified in these fields do not overlap with the ports you designate for user traffic. These values can only be changed by editing the corresponding Windows registry settings.

1. Select **Configure**.
2. Set the **Source Port Allocation Range** (default 20000-39999). This is the full range of port numbers that the TS Agent will allocate for user mapping. The port range you specify cannot overlap with the **System Source Port Allocation Range**.
3. (Optional) If there are ports/port ranges within the source port allocation that you do not want the TS Agent to allocate to user sessions, specify them as **Reserved Source Ports**. To include multiple ranges, use commas with no spaces, for example: `2000-3000,3500,4000-5000`.
4. Specify the number of ports to allocate to each individual user upon login to the terminal server in the **Port Allocation Start Size Per User** field (default 200).
5. Specify the **Port Allocation Maximum Size Per User**, which is the maximum number of ports the Terminal Server agent can allocate to an individual user.
6. Specify whether to continue processing traffic from the user if the user runs out of allocated ports. By default, the **Fail port binding when available ports are used up** is selected, which indicates that the application will fail to send traffic when all ports are used. To enable users to continue using applications when they run out of ports, clear this check box. Keep in mind that this traffic may not be identified with User-ID.

Step 5 Configure the firewalls to connect to the Terminal Server agent.

Complete the following steps on each firewall you want to connect to the Terminal Server agent to receive user mappings:

1. Select **Device > User Identification > Terminal Server Agents** and click **Add**.
2. Enter a **Name** for the Terminal Server agent.
3. Enter the IP address of the Windows **Host** on which the Terminal Server agent is installed.
4. Enter the **Port** number on which the agent will listen for user mapping requests. This value must match the value configured on the Terminal Server agent. By default, the port is set to 5009 on the firewall and on the agent. If you change it here, you must also change the **Listening Port** field on the Terminal Server agent **Configure** screen.
5. Make sure that the configuration is **Enabled** and then click **OK**.
6. **Commit** the changes.
7. Verify that the **Connected** status displays as connected (a green light).

| Install the Windows Terminal Server Agent (Continued) | |
|---|---|
| Step 6 Verify that the Terminal Server agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.<br> 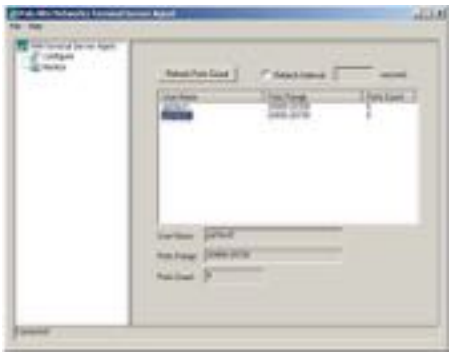 | 1. Launch the Terminal Server agent and verify that the firewalls can connect by making sure the **Connection Status** for each of Device in the Connection List is **Connected**.<br>2. To verify that the Terminal Server agent is successfully mapping port ranges to usernames, select **Monitoring** and make sure that the mapping table is populated. |

## Retrieve User Mappings from a Terminal Server Using the User-ID XML API

The User-ID XML API is a RESTful API that uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services.

To enable a non-Windows terminal server to send user mapping information directly to the firewall, create scripts that extract the user login and logout events and use them for input to the User-ID XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget and providing the firewall's API key for secure communication. Creating user mappings from multi-user systems such as terminal servers requires use of the following API messages:

- **<multiusersystem>**—Sets up the configuration for an XML API Multi-user System on the firewall. This message allows for definition of the terminal server IP address (this will be the source address for all users on that terminal server). In addition, the <multiusersystem> setup message specifies the range of source port numbers to allocate for user mapping and the number of ports to allocate to each individual user upon login (called the *block size*). If you want to use the default source port allocation range (1025-65534) and block size (200), you do not need to send a <multiusersystem> setup event to the firewall. Instead, the firewall will automatically generate the XML API Multi-user System configuration with the default settings upon receipt of the first user login event message.
- **<blockstart>**—Used with the <login> and <logout> messages to indicate the starting source port number allocated to the user. The firewall then uses the block size to determine the actual range of port numbers to map to the IP address and username in the login message. For example, if the <blockstart> value is 13200 and the block size configured for the multi-user system is 300, the actual source port range allocated to the user is 13200 through 13499. Each connection initiated by the user should use a unique source port number within the allocated range, enabling the firewall to identify the user based on its IP address-port-user mappings for enforcement of user- and group-based security policy rules. When a user exhausts all the ports allocated, the terminal server must send a new <login> message allocating a new port range for the user so that the firewall can update the IP address-port-user mapping. In addition, a single

 © Palo Alto Networks, Inc.

username can have multiple blocks of ports mapped simultaneously. When the firewall receives a <logout> message that includes a <blockstart> parameter, it removes the corresponding IP address-port-user mapping from its mapping table. When the firewall receives a <logout> message with a username and IP address, but no <blockstart>, it removes the user from its table. And, if the firewall receives a <logout> message with an IP address only, it removes the multi-user system and all mappings associated with it.

The XML files that the terminal server sends to the firewall can contain multiple message types and the messages do not need to be in any particular order within the file. However, upon receiving an XML file that contains multiple message types, the firewall will process them in the following order: multiusersystem requests first, followed by logins, then logouts.

The following workflow provides an example of how to use the User-ID XML API to send user mappings from a non-Windows terminal server to the firewall.

**Use the User-ID XML API to Map Non-Windows Terminal Services Users**

Step 1 — Generate the API key that will be used to authenticate the API communication between the firewall and the Terminal server. To generate the key you must provide login credentials for an administrative account; the API is available to all administrators (including role-based administrators with XML API privileges enabled).

Any special characters in the password must be URL/ percent-encoded.

From a browser, log in to the firewall. Then, to generate the API key for the firewall, open a new browser window and enter the following URL:

```
https://<Firewall-IPaddress>/api/?type=keygen&user=<username>&
password=<password>
```

Where <Firewall-IPaddress> is the IP address or FQDN of the firewall and <username> and <password> are the credentials for the administrative user account on the firewall. For example:

```
https://10.1.2.5/api/?type=keygen&user=admin&password=admin
```

The firewall responds with a message containing the key, for example:

```
<response status="success">
   <result>
      <key>k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg=</key>
   </result>
</response>
```

**Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)**

Step 2 (Optional) Generate a setup message that the terminal server will send to specify the port range and block size of ports per user that your terminal services agent uses.

If the terminal services agent does not send a setup message, the firewall will automatically create a terminal server agent configuration using the following default settings upon receipt of the first login message:

- Default port range: 1025 to 65534
- Per user block size: 200
- Maximum number of multi-user systems: 1000

The following shows a sample setup message:

```
<uid-message>
  <payload>
    <multiusersystem>
      <entry ip="10.1.1.23" startport="20000"
        endport="39999" blocksize="100">
    </multiusersystem>
  </payload>
  <type>update</type>
  <version>1.0</version>
</uid-message>
```

where `entry ip` specifies the IP address assigned to terminal server users, `startport` and `endport` specify the port range to use when assigning ports to individual users, and `blocksize` specifies the number of ports to assign to each user. The maximum blocksize is 4000 and each multi-user system can allocate a maximum of 1000 blocks.

If you define a custom blocksize and or port range, keep in mind that you must configure the values such that every port in the range gets allocated and that there are no gaps or unused ports. For example, if you set the port range to 1000-1499, you could set the block size to 100, but not to 200. This is because if you set it to 200, there would be unused ports at the end of the range.

Step 3 Create a script that will extract the login events and create the XML input file to send to the firewall.

Make sure the script enforces assignment of port number ranges at fixed boundaries with no port overlaps. For example, if the port range is 1000-1999 and the block size is 200, acceptable blockstart values would be 1000, 1200, 1400, 1600, or 1800. Blockstart values of 1001, 1300, or 1850 would be unacceptable because some of the port numbers in the range would be left unused.

The login event payload that the terminal server sends to the firewall can contain multiple login events.

The following shows the input file format for a user-ID XML login event:

```
<uid-message>
<payload>
<login>
<entry name="acme\jjaso" ip="10.1.1.23" blockstart="20000">
<entry name="acme\jparker" ip="10.1.1.23" blockstart="20100">
<entry name="acme\ccrisp" ip="10.1.1.23" blockstart="21000">
</login>
</payload>
<type>update</type>
<version>1.0</version>
</uid-message>
```

The firewall uses this information to populate its user mapping table. Based on the mappings extracted from the example above, if the firewall received a packet with a source address and port of 10.1.1.23:20101, it would map the request to user jparker for policy enforcement.

Each multi-user system can allocate a maximum of 1000 port blocks.

 © Palo Alto Networks, Inc.

**Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)**

**Step 4** Create a script that will extract the logout events and create the XML input file to send to the firewall.

Upon receipt of a `logout` event message with a `blockstart` parameter, the firewall removes the corresponding IP address-port-user mapping. If the `logout` message contains a username and IP address, but no `blockstart` parameter, the firewall removes all mappings for the user. If the `logout` message contains an IP address only, the firewall removes the multi-user system and all associated mappings.

The following shows the input file format for a User-ID XML logout event:

```
<uid-message>
  <payload>
     <logout>
        <entry name="acme\jjaso" ip="10.1.1.23"
blockstart="20000">
        <entry name="acme\ccrisp" ip="10.1.1.23">
        <entry ip="10.2.5.4">
     </logout>
  </payload>
  <type>update</type>
  <version>1.0</version>
</uid-message>
```

You can also clear the multiuser system entry from the firewall using the following CLI command: **`clear xml-api multiusersystem`**

**Step 5** Make sure that the scripts you create include a way to dynamically enforce that the port block range allocated using the XML API matches the actual source port assigned to the user on the terminal server and that the mapping is removed when the user logs out or the port allocation changes.

One way to do this would be to use netfilter NAT rules to hide user sessions behind the specific port ranges allocated via the XML API based on the uid. For example, to ensure that a user with the user ID jjaso is mapped to a source network address translation (SNAT) value of 10.1.1.23:20000-20099, the script you create should include the following:

```
[root@ts1 ~]# iptables -t nat -A POSTROUTING -m owner --uid-owner jjaso
-p tcp -j SNAT --to-source 10.1.1.23:20000-20099
```

Similarly, the scripts you create should also ensure that the IP table routing configuration dynamically removes the SNAT mapping when the user logs out or the port allocation changes:

```
[root@ts1 ~]# iptables -t nat -D POSTROUTING 1
```

**Step 6** Define how to package the XML input files containing the setup, login, and logout events into wget or cURL messages for transmission to the firewall.

To apply the files to the firewall using wget:

```
> wget --post file <filename>
"https://<Firewall-IPaddress>/api/?type=user-id&key=<key>&file-name=<inp
ut_filename.xml>&client=wget&vsys=<VSYS_name>"
```

For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key `k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg` using wget would look as follows:

```
> wget --post file login.xml
"https://10.2.5.11/api/?type=user-id&key=k7J335J6hI7nBxIqyfa62sZugWx
7ot%2BgzEA9UOnlZRg&file-name=login.xml&client=wget&vsys=vsys1"
```

To apply the file to the firewall using cURL:

```
> curl --form file=@<filename>
https://<Firewall-IPaddress>/api/?type=user-id&key=<key>&vsys=<VSYS_name
>
```

For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key `k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg` using cURL would look as follows:

```
> curl --form file@login.xml
"https://10.2.5.11/api/?type=user-id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%
2BgzEA9UOnlZRg&vsys=vsys1"
```

| Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued) | |
|---|---|
| Step 7 Verify that the firewall is successfully receiving login events from the terminal servers. | Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:<br>To verify if the terminal server is connecting to the firewall over XML:<br>`admin@PA-5050> show user xml-api multiusersystem`<br>`Host Vsys Users Blocks`<br>`----------`<br>`10.5.204.43 vsys1 5 2`<br>To verify that the firewall is receiving mappings from a terminal server over XML:<br>`admin@PA-5050> show user ip-port-user-mapping all`<br>`Global max host index 1, host hash count 1`<br>`XML API Multi-user System 10.5.204.43`<br>`Vsys 1, Flag 3`<br>`Port range: 20000 - 39999`<br>`Port size: start 200; max 2000`<br>`Block count 100, port count 20000`<br>`20000-20199: acme\administrator`<br>`Total host: 1` |

 © Palo Alto Networks, Inc.

## Send User Mappings to User-ID Using the XML API

Although the User-ID functionality provides many out-of-the box methods for obtaining user mapping information, you may have some applications or devices that capture user information that cannot be natively integrated with User-ID. In this case you can use the User-ID XML API to create custom scripts that allow you to leverage existing user data and send it to the User-ID agent or directly to the firewall.

The User-ID XML API is a RESTful API that uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services. To leverage user data from an existing system—such as a custom application developed internally or another device that is not supported by one of the existing user mapping mechanisms—you can create custom scripts to extract the data and send it to the firewall or the User-ID agent using the XML API.

To enable an external system to send user mapping information to the User-ID agent or directly to the firewall, you can create scripts that extract the user login and logout events and use them for input to the User-ID XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget using the firewall's API key for secure communication. For more details, refer to the *PAN-OS XML API Usage Guide*.

# Configure a Firewall to Share User Mapping Data with Other Firewalls

Because policy is local to each firewall, each firewall needs current user mapping and group mapping information to accurately enforce security policy by user and group. However, you can configure one firewall to collect all the mapping information and act as a User-ID agent to share that information with other firewalls. The redistribution firewall can share only the information it collects using local methods (for example, the PAN-OS integrated User-ID agent or Captive Portal), not the information collected from the Windows-based User-ID and Terminal Services agents. You configure the receiving firewalls to retrieve the mapping information from the redistribution firewall; they don't need to communicate directly with domain servers.

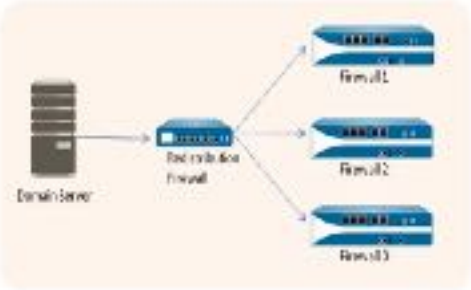

The following procedure describes how to set up redistribution of User-ID information.

| **Configure a Firewall to Redistribute User Mappings** | | |
|---|---|---|
| Step 1 | Configure the redistribution firewall.<br><br>User-ID configurations apply to a single virtual system only. To redistribute User-ID mappings from multiple virtual systems, you must configure the user mapping settings on each virtual system separately, using a unique pre-shared key in each configuration. | 1. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User-ID Agent Setup section.<br>2. Select **Redistribution**.<br>3. Enter a **Collector Name**.<br>4. Enter and confirm the **Pre-Shared Key** that will enable other firewalls to connect to this firewall to retrieve user mapping information.<br>5. Click **OK** to save the redistribution configuration. |
| Step 2 | Create an interface management profile that enables the User-ID service and attach it to the interface that the other firewalls will connect to in order to retrieve user mappings. | 1. Select **Network > Network Profiles > Interface Mgmt** and click **Add**.<br>2. Enter a **Name** for the profile and then select the Permitted Services. At a minimum, select **User-ID Service** and **HTTPS**.<br>3. Click **OK** to save the profile.<br>4. Select **Network > Interfaces > Ethernet** and select the interface you plan to use for redistribution.<br>5. On the **Advanced > Other Info** tab, select the **Management Profile** you just created.<br>6. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

| Configure a Firewall to Redistribute User Mappings (Continued) | | |
|---|---|---|
| Step 3 | Configure the other firewalls to retrieve user mappings from the redistribution firewall.<br><br>If the redistribution firewall has multiple virtual systems configured for redistribution, make sure you are using the pre-shared key that corresponds to the virtual system from which you want this firewall to retrieve User-ID mappings. | Perform the following steps on each firewall that you want to be able to retrieve user mappings:<br>1. Select **Device > User Identification > User-ID Agents**.<br>2. Click **Add** and enter a User-ID agent **Name** for the redistribution firewall.<br>3. Enter the hostname or IP address of the firewall interface that you configured for redistribution in the **Host** field.<br>4. Enter 5007 as the **Port** number on which the redistribution firewall will listen for User-ID requests.<br>5. Enter the **Collector Name** that you specified in the redistribution firewall configuration (Step 1-3).<br>6. Enter and confirm the **Collector Pre-Shared Key**. The key value you enter here must match the value configured on the redistribution firewall (Step 1-4).<br>7. (Optional) If you are using the redistribution firewall to retrieve group mappings in addition to user mappings, select the **Use as LDAP Proxy** check box.<br>8. (Optional) If you are using the redistribution firewall for Captive Portal authentication, select the **Use for NTLM Authentication** check box.<br>9. Make sure that the configuration is **Enabled** and then click **OK**.<br>10. **Commit** the changes. |
| Step 4 | Verify the configuration. | On the **User-ID Agents** tab, verify that the redistribution firewall entry you just added shows a green icon in the **Connected** column.<br><br>If a red icon appears, check traffic logs (**Monitor > Logs > Traffic**) to identify the issue. You can also check to see if any user mapping data has been received by running the following operational commands from the CLI:<br><br>`show user ip-user-mapping` (to view user mapping information on the dataplane)<br><br>`show user ip-user-mapping-mp` (to view mappings on the management plane). |

# Enable User- and Group-Based Policy

In order to enable security policy based on user and/or group, you must enable User-ID for each zone that contains users you want to identify. You can then define policies that allow or deny traffic based on user name or group membership. Additionally, you can create Captive Portal policies to enable identification for IP addresses that do not yet have any user data associated with them.

| Enable User- and Group-Based Policy | |
|---|---|
| Step 1 Enable User-ID on the source zones that contain the users that will send requests that require user-based access controls. | 1. Select **Network > Zones**.<br>2. Click the **Name** of the zone in which you want to enable User-ID to open the Zone dialog.<br>3. Select the **Enable User Identification** check box and then click **OK**.<br> 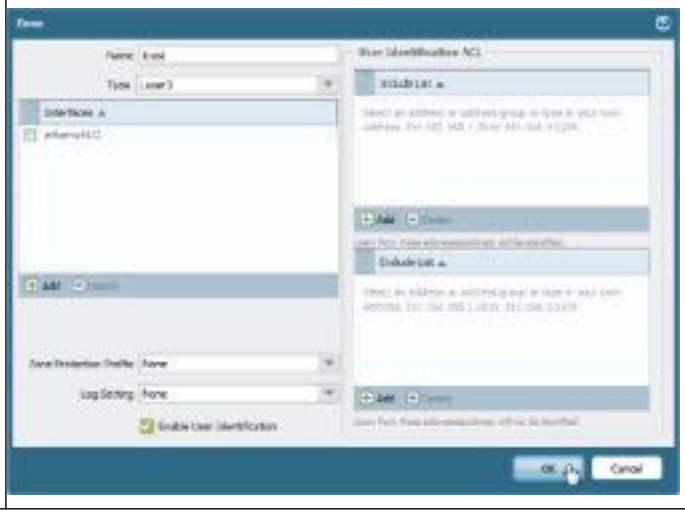 |

 © Palo Alto Networks, Inc.

| Enable User- and Group-Based Policy (Continued) | |
|---|---|
| Step 2 Create security policies based on user and/or group.<br> As a best practice, create policies based on group rather than user whenever possible. This prevents you from having to continually update your policies (which requires a commit) whenever your user base changes. | 1. After configuring User-ID, you will be able to choose a user name or group name when defining the source or destination of a security rule:<br>a. Select **Policies > Security** and click **Add** to create a new policy or click on an existing policy rule name to open the Security Policy Rule dialog.<br>b. Specify which users and/or groups to match in the policy in one of the following ways:<br>– If you want to specify specific users/groups as matching criteria, select the **User** tab and click the **Add** button in the Source User section to display a list of users and groups discovered by the firewall group mapping function. Select the users and/or groups to add to the policy.<br>– If you want the policy to match any user who has or has not successfully authenticated and you don't need to know the specific user or group name, select **known-user** or **unknown** from the drop-down above the **Source User** list.<br>2. Configure the rest of the policy as appropriate and then click **OK** to save it. For details on other fields in the security policy, see Set Up Basic Security Policies. |

**Enable User- and Group-Based Policy (Continued)**

| | |
|---|---|
| Step 3 Create your Captive Portal policies. | 1. Select **Policies > Captive Portal**.<br>2. Click **Add** and enter a **Name** for the policy.<br>3. Define the matching criteria for the rule by completing the **Source**, **Destination**, and **Service/URL Category** tabs as appropriate to match the traffic you want to authenticate. The matching criteria on these tabs is the same as the criteria you define when creating a security policy. See Set Up Basic Security Policies for details.<br>4. Define the **Action** to take on traffic that matches the rule:<br>• **no-captive-portal**—Allow traffic to pass without presenting a Captive Portal page for authentication.<br>• **web-form**—Present a Captive Portal page for the user to explicitly enter authentication credentials or use client certificate authentication.<br>• **browser-challenge**—Open an NTLM authentication request to the web browser of the user. The web browser will respond using the current login credentials of the user. If the login credentials are not available, the user will be prompted to supply them. If you select this action, you must enable NT LAN Manager (NTLM) authentication when you Configure Captive Portal. If NTLM authentication fails, the firewall falls back to **web-form** authentication.<br>5. Click **OK**.<br>The following example shows a Captive Portal policy that instructs the firewall to present a web form to authenticate unknown users who send HTTP requests from the trust zone to the untrust zone. |
| Step 4 Save your policy settings. | Click **Commit**. |

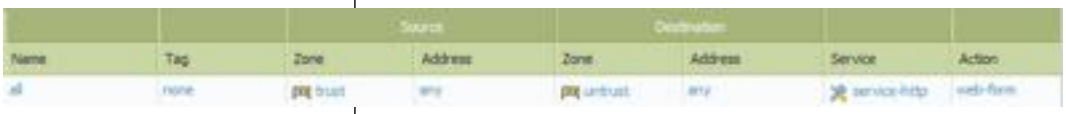


 © Palo Alto Networks, Inc.

# Verify the User-ID Configuration

After you configure group mapping and user mapping and enable User-ID on your security policies and Captive Portal policies, you should verify that it is working properly.

| Verify the User-ID Configuration | |
|---|---|
| Step 1 Verify that group mapping is working. | From the CLI, enter the following command:<br>**show user group-mapping statistics** |
| Step 2 Verify that user mapping is working. | If you are using the on-device User-ID agent, you can verify this from the CLI using the following command:<br>**show user ip-user-mapping-mp all**<br>`IP Vsys From User Timeout (sec)`<br>`----------`<br>`192.168.201.1 vsys1 UIA acme\george 210`<br>`192.168.201.11 vsys1 UIA acme\duane 210`<br>`192.168.201.50 vsys1 UIA acme\betsy 210`<br>`192.168.201.10 vsys1 UIA acme\administrator 210`<br>`192.168.201.100 vsys1 AD acme\administrator 748`<br>`Total: 5 users`<br>`*: WMI probe succeeded` |
| Step 3 Test your security policy. | • From a machine in the zone where User-ID is enabled, attempt to access sites and applications to test the rules you defined in your policy and ensure that traffic is allowed and denied as expected.<br>• You can also use the `test security-policy-match` command to determine whether the policy is configured correctly. For example, suppose you have a rule that blocks user duane from playing World of Warcraft; you could test the policy as follows:<br>`test security-policy-match application worldofwarcraft source-user acme\duane source any destination any destination-port any protocol 6`<br>`"deny worldofwarcraft" {`<br>`from corporate;`<br>`source any;`<br>`source-region any;`<br>`to internet;`<br>`destination any;`<br>`destination-region any;`<br>`user acme\duane;`<br>`category any;`<br>`application/service worldofwarcraft;`<br>`action deny;`<br>`terminal no;`<br>`}` |

**Verify the User-ID Configuration (Continued)**

Step 4 Test your Captive Portal configuration.

1. From the same zone, go to a machine that is not a member of your directory, such as a Mac OS system, and try to ping to a system external to the zone. The ping should work without requiring authentication.
2. From the same machine, open a browser and navigate to a web site in a destination zone that matches a Captive Portal policy you defined. You should see the Captive Portal web form.

3. Log in using the correct credentials and confirm that you are redirected to the requested page.
4. You can also test your Captive Portal policy using the `test cp-policy-match` command as follows:

```
test cp-policy-match from corporate to internet
source 192.168.201.10 destination 8.8.8.8
Matched rule: 'captive portal' action: web-form
```

Step 5 Verify that user names are displayed in the log files (**Monitor > Logs**).

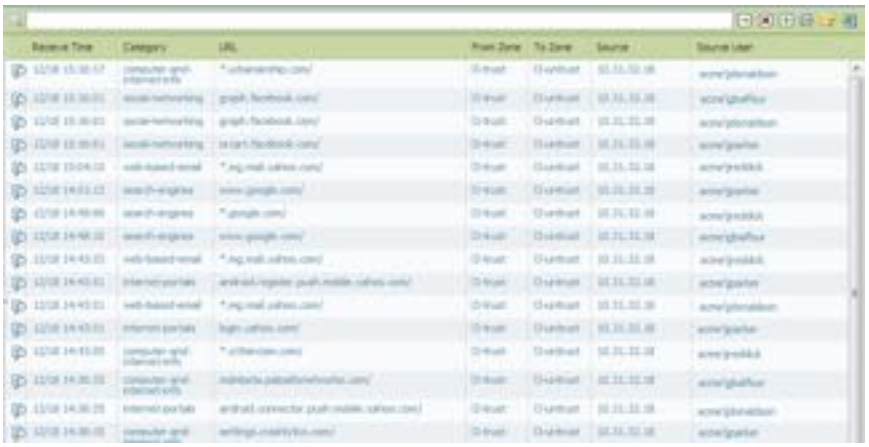

 © Palo Alto Networks, Inc.

**Verify the User-ID Configuration (Continued)**

Step 6 Verify that user names are displayed in reports (**Monitor > Reports**). For example, when drilling down into the denied applications report, you should see a list of the users who attempted to access the applications as in the following example.

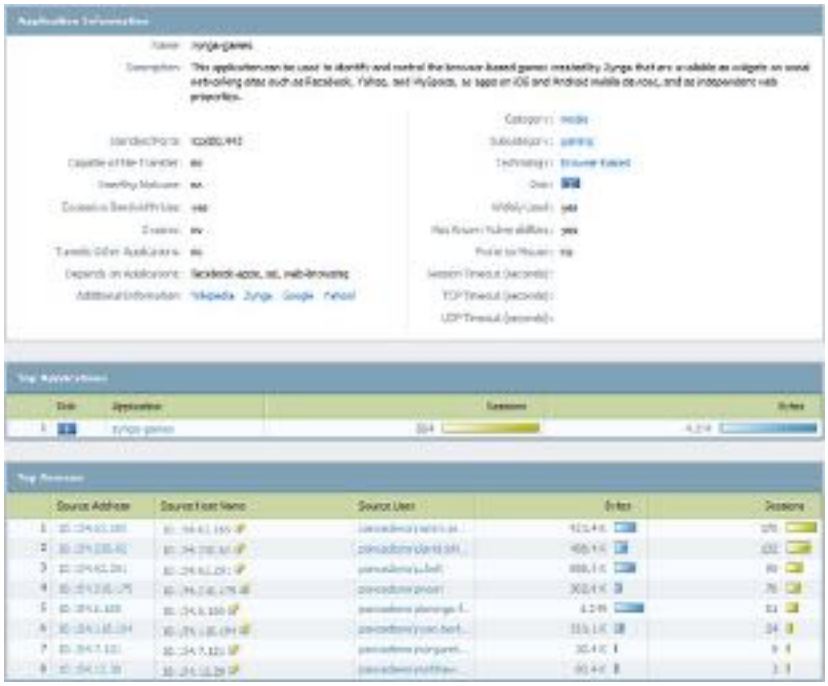

 © Palo Alto Networks, Inc.

# App-ID

To safely enable applications on your network, the Palo Alto Networks next-generation firewalls provide both an application and web perspective—App-ID and URL Filtering—to protect against a full spectrum of legal, regulatory, productivity, and resource utilization risks.

App-ID enables visibility into the applications on the network, so you can learn how they work and understand their behavioral characteristics and their relative risk. This application knowledge allows you to create and enforce security policies to enable, inspect, and shape desired applications and block undesired applications. When you define policies to begin allowing traffic, App-ID begins to classify traffic without any additional configuration.


- App-ID Overview
- Manage Custom or Unknown Applications
- Best Practices for Using App-ID in Policy
- Applications with Implicit Support
- Application Level Gateways
- Disable the SIP Application-level Gateway (ALG)

# App-ID Overview

App-ID, a patented traffic classification system only available in Palo Alto Networks firewalls, determines what the application is irrespective of port, protocol, encryption (SSH or SSL) or any other evasive tactic used by the application. It applies multiple classification mechanisms—application signatures, application protocol decoding, and heuristics—to your network traffic stream to accurately identify applications.

Here's how App-ID identifies applications traversing your network:

- Traffic is matched against policy to check whether it is allowed on the network.
- Signatures are then applied to allowed traffic to identify the application based on unique application properties and related transaction characteristics. The signature also determines if the application is being used on its default port or it is using a non-standard port. If the traffic is allowed by policy, the traffic is then scanned for threats and further analyzed for identifying the application more granularly.
- If App-ID determines that encryption (SSL or SSH) is in use, and a decryption policy is in place, the session is decrypted and application signatures are applied again on the decrypted flow.
- Decoders for known protocols are then used to apply additional context-based signatures to detect other applications that may be tunneling inside of the protocol (e.g., Yahoo! Instant Messenger used across HTTP). Decoders validate that the traffic conforms to the protocol specification and provide support for NAT traversal and opening dynamic pinholes for applications such as SIP and FTP.
- For applications that are particularly evasive and cannot be identified through advanced signature and protocol analysis, heuristics or behavioral analysis may be used to determine the identity of the application.

When the application is identified, the policy check determines how to treat the application, for example—block, or allow and scan for threats, inspect for unauthorized file transfer and data patterns, or shape using QoS.

 © Palo Alto Networks, Inc.

# Manage Custom or Unknown Applications

Palo Alto Networks provides weekly App-ID updates to identify new applications. By default, App-ID is always enabled on the firewall, and you don't need to enable a series of signatures to identify well-known applications. Typically, the only applications that are classified as *unknown* traffic—*tcp, udp or non-syn-tcp*—in the ACC and the traffic logs are commercially available applications that have not yet been added to App-ID, internal or custom applications on your network, or potential threats.

On occasion, the firewall may report an application as unknown for the following reasons:

- Incomplete data—A handshake took place, but no data packets were sent prior to the timeout.
- Insufficient data—A handshake took place followed by one or more data packets; however, not enough data packets were exchanged to identify the application.

The following choices are available to handle unknown applications:

- Create security policies to control unknown applications by unknown TCP, unknown UDP or by a combination of source zone, destination zone, and IP addresses.
- Request an App-ID from Palo Alto Networks—If you would like to inspect and control the applications that traverse your network, for any unknown traffic, you can record a packet capture. If the packet capture reveals that the application is a commercial application, you can submit this packet capture to Palo Alto Networks for App-ID development. If it is an internal application, you can create a custom App-ID and/or define an application override policy.
- Create a custom App-ID with a signature and attach it to a security policy, or create a custom App-ID and define an application override policy—A custom App-ID allows you to customize the definition of the internal application—its characteristics, category and sub-category, risk, port, timeout—and exercise granular policy control in order to minimize the range of unidentified traffic on your network. Creating a custom App-ID also allows you to correctly identify the application in the ACC and traffic logs and is useful in auditing/reporting on the applications on your network. For a custom application you can specify a signature and a pattern that uniquely identifies the application and attach it to a security policy that allows or denies the application.

In order to collect the right data to create a custom application signature, you'll need a good understanding of packet captures and how datagrams are formed. If the signature is created too broadly, you might inadvertently include other similar traffic; if it is defined too narrowly, the traffic will evade detection if it does not strictly match the pattern.

Custom App-IDs are stored in a separate database on the firewall and this database is not impacted by the weekly App-ID updates.

The supported application protocol decoders that enable the firewall to detect applications that may be tunneling inside of the protocol include the following as of content update 424: HTTP, HTTPS, DNS, FTP, IMAP SMTP, Telnet, IRC (Internet Relay Chat), Oracle, RTMP, RTSP, SSH, GNU-Debugger, GIOP (Global Inter-ORB Protocol), Microsoft RPC, Microsoft SMB (also known as CIFS). Additionally, with the 4.0 release of PAN-OS, this custom App-ID capability has been expanded to include unknown TCP and unknown UDP.

Alternatively, if you would like the firewall to process the custom application using fast path (Layer-4 inspection instead of using App-ID for Layer-7 inspection), you can reference the custom App-ID in an application override policy. An application override with a custom application will prevent the session from being processed by the App-ID engine, which is a Layer-7 inspection. Instead it forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4, and thereby saves application processing time.

For example, if you build a custom application that triggers on a host header *www.mywebsite.com*, the packets are first identified as *web-browsing* and then are matched as your custom application (whose parent application is web-browsing). Because the parent application is web-browsing, the custom application is inspected at Layer-7 and scanned for content and vulnerabilities.

If you define an application override, the firewall stops processing at Layer-4. The custom application name is assigned to the session to help identify it in the logs, and the traffic is not scanned for threats.

For more details, refer to the following articles:

- Identifying Unknown Applications
- Video: How to Configure a Custom App-ID
- Custom Application Signatures

 © Palo Alto Networks, Inc.

# Best Practices for Using App-ID in Policy

- Review the ACC for the list of applications on your network, and determine which applications to allow or block. If you are migrating from a firewall where you defined port-based rules, to get a list of applications that run on a given port, search by the port number in the application browser (**Objects > Applications**) on the Palo Alto Networks firewall or in Applipedia.
- Use **application-default** for the **Service**. The firewall compares the port used with the list of default ports for that application. If the port used is not a default port for the application, the firewall drops the session and logs the message **appid policy lookup deny**. If you have an application that is accessed on many ports and you would like to limit the ports on which the application is used, specify it in **Service**/**Service Group** objects in policies.
- Use application filters to dynamically include new applications in existing policy rules. See an example.

# Applications with Implicit Support

When creating a policy to allow specific applications, you must also be sure that you are allowing any other applications on which the application depends. In many cases, you do not have to explicitly allow access to the dependent applications in order for the traffic to flow because the firewall is able to determine the dependencies and allow them implicitly. This implicit support also applies to custom applications that are based on HTTP, SSL, MS-RPC, or RTSP. Applications for which the firewall cannot determine dependent applications on time will require that you explicitly allow the dependent applications when defining your policies. You can determine application dependencies in Applipedia.

The following table lists the applications for which the firewall has implicit support (as of Content Update 484).

**Table: Applications with Implicit Support**

| Application | Implicitly Supports |
|---|---|
| 360-safeguard-update | http |
| apple-update | http |
| apt-get | http |
| as2 | http |
| avg-update | http |
| avira-antivir-update | http, ssl |
| blokus | rtmp |
| bugzilla | http |
| clubcooee | http |
| corba | http |
| cubby | http, ssl |
| dropbox | ssl |
| esignal | http |
| evernote | http, ssl |
| ezhelp | http |
| facebook | http, ssl |
| facebook-chat | jabber |
| facebook-social-plugin | http |
| fastviewer | http, ssl |
| forticlient-update | http |
| good-for-enterprise | http, ssl |
| google-cloud-print | http, ssl, jabber |
| google-desktop | http |

 © Palo Alto Networks, Inc.

| Application | Implicitly Supports |
|---|---|
| google-drive-web | http |
| google-talk | jabber |
| google-update | http |
| gotomypc-desktop-sharing | citrix-jedi |
| gotomypc-file-transfer | citrix-jedi |
| gotomypc-printing | citrix-jedi |
| hipchat | http |
| iheartradio | ssl, http, rtmp |
| infront | http |
| instagram | http, ssl |
| issuu | http, ssl |
| java-update | http |
| jepptech-updates | http |
| kerberos | rpc |
| kik | http, ssl |
| lastpass | http, ssl |
| logmein | http, ssl |
| mcafee-update | http |
| megaupload | http |
| metatrader | http |
| mocha-rdp | t_120 |
| mount | rpc |
| ms-frs | msrpc |
| ms-rdp | t_120 |
| ms-scheduler | msrpc |
| ms-service-controller | msrpc |
| nfs | rpc |
| oovoo | http, ssl |
| paloalto-updates | ssl |
| panos-global-protect | http |
| panos-web-interface | http |
| pastebin | http |

© Palo Alto Networks, Inc.

| Application | Implicitly Supports |
|---|---|
| pastebin-posting | http |
| pinterest | http, ssl |
| portmapper | rpc |
| prezi | http, ssl |
| rdp2tcp | t_120 |
| renren-im | jabber |
| roboform | http, ssl |
| salesforce | http |
| stumbleupon | http |
| supremo | http |
| symantec-av-update | http |
| trendmicro | http |
| trillian | http, ssl |
| twitter | http |
| whatsapp | http, ssl |
| xm-radio | rtsp |


© Palo Alto Networks, Inc.

# Application Level Gateways

The Palo Alto Networks firewall does not classify traffic by port and protocol; instead it identifies the application based on its unique properties and transaction characteristics using the App-ID technology. Some applications, however, require the firewall to dynamically open *pinholes* to establish the connection, determine the parameters for the session and negotiate the ports that will be used for the transfer of data; these applications use the application-layer payload to communicate the dynamic TCP or UDP ports on which the application opens data connections. For such applications, the firewall serves as an Application Level Gateway (ALG), and it opens a pinhole for a limited time and for exclusively transferring data or control traffic. The firewall also performs a NAT rewrite of the payload when necessary.

As of Content Release version 464, the Palo Alto Networks firewall provides NAT ALG support for the following protocols: FTP, H.225, H.248, MGCP, MySQL, Oracle/SQLNet/TNS, RPC, RTSP, SCCP, SIP, and UNIStim.

When the firewall serves as an ALG for the Session Initiation Protocol (SIP), by default it performs NAT on the payload and opens dynamic pinholes for media ports. In some cases, depending on the SIP applications in use in your environment, the SIP endpoints have NAT intelligence embedded in their clients. In such cases, you might need to disable the SIP ALG functionality to prevent the firewall from modifying the signaling sessions. When SIP ALG is disabled, if App-ID determines that a session is SIP, the payload is not translated and dynamic pinholes are not opened. See Disable the SIP Application-level Gateway (ALG).

© Palo Alto Networks, Inc.

# Disable the SIP Application-level Gateway (ALG)

The Palo Alto Networks firewall uses the Session Initiation Protocol (SIP) application-level gateway (ALG) to open dynamic pinholes in the firewall where NAT is enabled. However, some applications—such as VoIP—have NAT intelligence embedded in the client application. In these cases, the SIP ALG on the firewall can interfere with the signaling sessions and cause the client application to stop working.

One solution to this problem is to define an Application Override Policy for SIP, but using this approach disables the App-ID and threat detection functionality. A better approach is to disable the SIP ALG, which does not disable App-ID or threat detection.

The following procedure describes how to disable the SIP ALG.

**Disable the SIP ALG**

1. Select **Objects > Applications**.
2. Select the **sip** application.
   You can type **sip** in the **Search** box to help find the sip application.
3. Select **Customize...** for **ALG** in the Options section of the Application dialog box.

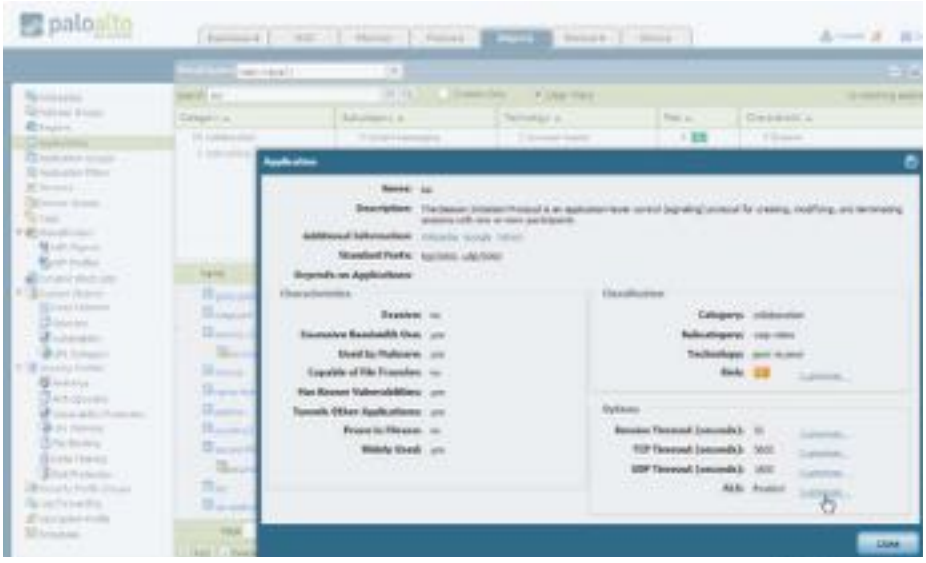

4. Select the **Disable ALG** check box in the Application - sip dialog box and click **OK**.

5. **Close** the Application dialog box and **Commit** the change.


© Palo Alto Networks, Inc.

# Threat Prevention

The Palo Alto Networks next-generation firewall protects and defends your network from commodity threats and advanced persistent threats (APTs). The firewall's multi-pronged detection mechanisms include a signature-based (IPS/Command and Control/Antivirus) approach, heuristics-based (bot detection) approach, sandbox-based (WildFire) approach, and Layer 7 protocol analysis-based (App-ID) approach.

Commodity threats are exploits that are less sophisticated and more easily detected and prevented using a combination of the antivirus, anti-spyware, vulnerability protection and the URL filtering/Application identification capabilities on the firewall.

Advanced threats are perpetuated by organized cyber criminals or malicious groups that use sophisticated attack vectors to target your network, most commonly for intellectual property theft and financial data theft. These threats are more evasive and require intelligent monitoring mechanisms for detailed host and network forensics on malware. The Palo Alto Networks next-generation firewall in conjunction with WildFire and Panorama provides a comprehensive solution that intercepts and break the attack chain and provides visibility to prevent security infringement on your network—including mobile and virtualized—infrastructure.


- Set Up Security Profiles and Policies
- Prevent Brute Force Attacks
- Customize the Action and Trigger Conditions for a Brute Force Signature
- Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions
- Enable Passive DNS Collection for Improved Threat Intelligence
- Use DNS Queries to Identify Infected Hosts on the Network
- Content Delivery Network Infrastructure for Dynamic Updates
- Threat Prevention Resources


© Palo Alto Networks, Inc.

# Set Up Security Profiles and Policies

The following sections provide basic threat prevention configuration examples:

- Set Up Antivirus, Anti-spyware, and Vulnerability Protection
- Set Up Data Filtering
- Set Up File Blocking

For information on controlling web access as part of your threat prevention strategy, see URL Filtering.

 © Palo Alto Networks, Inc.

## Set Up Antivirus, Anti-spyware, and Vulnerability Protection

The following describes the steps needed to set up the default Antivirus, Anti-spyware, and Vulnerability Protection Security Profiles.

All Anti-spyware and Vulnerability Protection signatures have a default action defined by Palo Alto Networks. You can view the default action by navigating to **Objects > Security Profiles > Anti-Spyware** or **Objects > Security Profiles >Vulnerability Protection** and then selecting a profile. Click the **Exceptions** tab and then click **Show all signatures** and you will see a list of the signatures with the default action in the Action column. To change the default action, you must create a new profile and then create rules with a non-default action, and/or add individual signature exceptions to **Exceptions** in the profile.

| **Set up Antivirus/Anti-spyware/Vulnerability Protection** | | |
|---|---|---|
| Step 1 | Verify that you have a Threat Prevention license. | • The Threat Prevention license bundles the Antivirus, Anti-spyware, and the Vulnerability Protection features in one license.<br>• Select **Device > Licenses** to verify that the **Threat Prevention** license is installed and check the expiration date. |
| Step 2 | Download the latest antivirus threat signatures. | 1. Select **Device > Dynamic Updates** and click **Check Now** at the bottom of the page to retrieve the latest signatures.<br>In the **Actions** column, click **Download** to install the latest Antivirus and Applications and Threats signatures. |
| Step 3 | Schedule signature updates. | 1. From **Device > Dynamic Updates**, click the text to the right of **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.<br>2. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select **Download Only**, you would need to manually go in and click the **Install** link in the **Action** column to install the signature. When you click **OK**, the update is scheduled. No commit is required.<br>3. (Optional) You can also enter the number of hours in the **Threshold** field to indicate the minimum age of a signature before a download will occur. For example, if you entered **10**, the signature must be at least 10 hours old before it will be downloaded, regardless of the schedule.<br>4. In an HA configuration, you can also click the **Sync To Peer** option to synchronize the content update with the HA peer after download/install. This will not push the schedule settings to the peer device, you need to configure the schedule on each device. |

**Set up Antivirus/Anti-spyware/Vulnerability Protection (Continued)**

### Best Practices for Antivirus Schedules

The general recommendation for antivirus signature update schedules is to perform a **download-and-install** on a daily basis for antivirus and weekly for applications and vulnerabilities.

### Recommendations for HA Configurations:

- **Active/Passive HA**—If the MGT port is used for antivirus signature downloads, you should configure a schedule on both devices and both devices will download/install independently. If you are using a data port for downloads, the passive device will not perform downloads while it is in the passive state. In this case you would set a schedule on both devices and then select the **Sync To Peer** option. This will ensure that whichever device is active, the updates will occur and will then push to the passive device.
- **Active/Active HA**—If the MGT port is used for antivirus signature downloads on both devices, then schedule the download/install on both devices, but do not select the **Sync To Peer** option. If you are using a data port, schedule the signature downloads on both devices and select **Sync To Peer**. This will ensure that if one device in the active/active configuration goes into the active-secondary state, the active device will download/install the signature and will then push it to the active-secondary device.

| | |
|---|---|
| Step 4 Attach the security profiles to a security policy. | 1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab.<br>2. In **Profile Settings**, click the drop-down next to each security profile you would like to enable. In this example we choose default for **Antivirus, Vulnerability Protection,** and **Anti-spyware**.<br>If no security profiles have been previously defined, select **Profiles** from the **Profile Type** drop-down. You will then see the list of options to select the security profiles.<br> 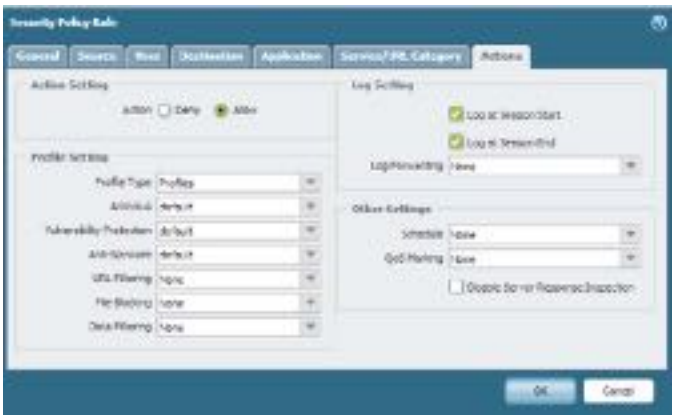 |
| Step 5 Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

## Set Up Data Filtering

The following describes the steps needed to configure a data filtering profile that will detect Social Security Numbers and a custom pattern identified in .doc and .docx documents.

| Data Filtering Configuration Example | | |
|---|---|---|
| Step 1 | Create a Data Filtering security profile. | 1. Select **Objects > Security Profiles > Data Filtering** and click **Add**.<br>2. Enter a **Name** and a **Description** for the profile. In this example the name is *DF_Profile1* with the description *Detect Social Security Numbers*.<br>3. (Optional) If you want to collect data that is blocked by the filter, select the **Data Capture** check box.<br> You must set a password as described in Step 2 if you are using the data capture feature. |
| Step 2 | (Optional) Secure access to the data filtering logs to prevent other administrators from viewing sensitive data.<br><br>When you enable this option, you will be prompted for the password when you view logs in **Monitor > Logs > Data Filtering**. | 1. Select **Device > Setup > Content-ID**.<br>2. Click **Manage Data Protection** in the Content-ID Features section.<br>3. Set the password that will be required to view the data filtering logs. |

**Data Filtering Configuration Example (Continued)**

Step 3 Define the data pattern that will be used in the Data Filtering Profile.

In this example, we will use the keyword **confidential** and will set the option to search for SSN numbers with dashes (Example - 987-654-4320).

It is helpful to set the appropriate thresholds and define keywords within documents to reduce false positives.

1. From the Data Filtering Profile page click **Add** and select **New** from the **Data Pattern** drop-down. You can also configure data patterns from **Objects > Custom Signatures > Data Patterns**.
2. For this example, name the Data Pattern signature **Detect SS Numbers** and add the description **Data Pattern to detect Social Security numbers**.
3. In the **Weight** section for **SSN#** enter **3**. See Weight and Threshold Values for more details.

4. (Optional) You can also set **Custom Patterns** that will be subject to this profile. In this case, you specify a pattern in the custom patterns **Regex** field and set a weight. You can add multiple match expressions to the same data pattern profile. In this example, we will create a **Custom Pattern** named **SSN_Custom** with a custom pattern of **confidential** (the pattern is case sensitive) and use a weight of **20**. The reason we use the term **confidential** in this example is because we know that our social security Word docs contain this term, so we define that specifically.

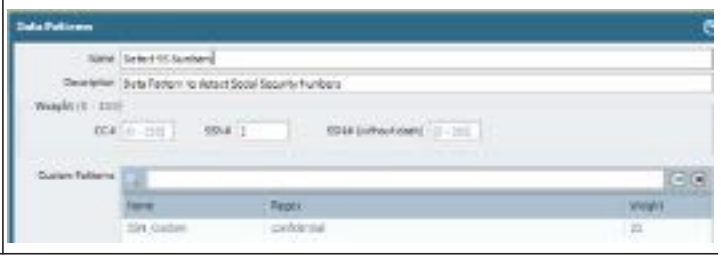

 © Palo Alto Networks, Inc.

| Data Filtering Configuration Example (Continued) | | |
|---|---|---|
| Step 4 | Specify which applications to filter and set the file types. | 1. Set **Applications** to **Any.** This will detect any supported application such as: web-browsing, FTP, or SMTP. If you want to narrow down the application, you can select it from the list. For applications such as Microsoft Outlook Web App that uses SSL, you will need to enable decryption. Also make sure you understand the naming for each application. For example, Outlook Web App, which is the Microsoft name for this application is identified as the application **outlook-web** in the PAN-OS list of applications. You can check the logs for a given application to identify the name defined in PAN-OS.<br>2. Set **File Types** to **doc** and **docx** to only scan doc and docx files. |
| Step 5 | Specify the direction of traffic to filter and the threshold values. | 1. Set the **Direction** to **Both**. Files that are uploaded or downloaded will be scanned.<br>2. Set the **Alert Threshold** to **35.** In this case, an alert will be triggered if 5 instances of Social Security Numbers exist and 1 instance of the term **confidential** exists. The formula is 5 SSN instances with a weight of 3 = 15 plus 1 instance of the term confidential with a weight of 20 = 35.<br>3. Set the **Block Threshold** to **50**. The file will be blocked if the threshold of 50 instances of a SSN and/or the term confidential exists in the file. In this case, if the doc contained 1 instance of the word **confidential** with a weight of 20 that equals 20 toward the threshold, and the doc has 15 Social Security Numbers with a weight of 3 that equals 45. Add 20 and 45 and you have 65, which will exceed the block threshold of 50. |

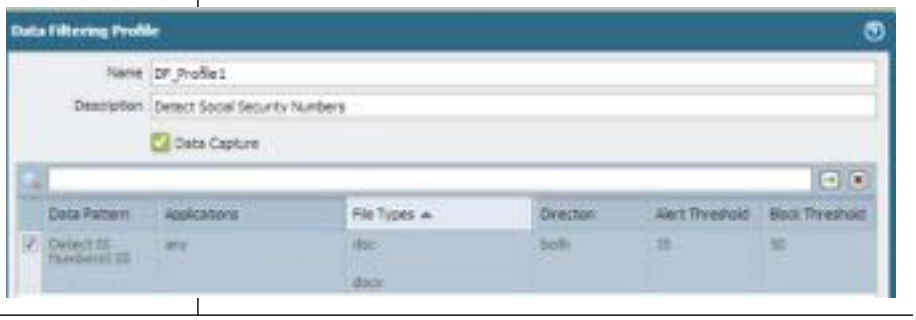

**Data Filtering Configuration Example (Continued)**

| | | |
|---|---|---|
| Step 6 | Attach the Data Filtering profile to the security rule. | 1. Select **Policies > Security** and select the security policy rule to which to apply the profile.<br>2. Click the security policy rule to modify it and then click the **Actions** tab. In the **Data Filtering** drop-down, select the new data filtering profile you created and then click **OK** to save. In this example, the data filtering rule name is **DF_Profile1**.<br> 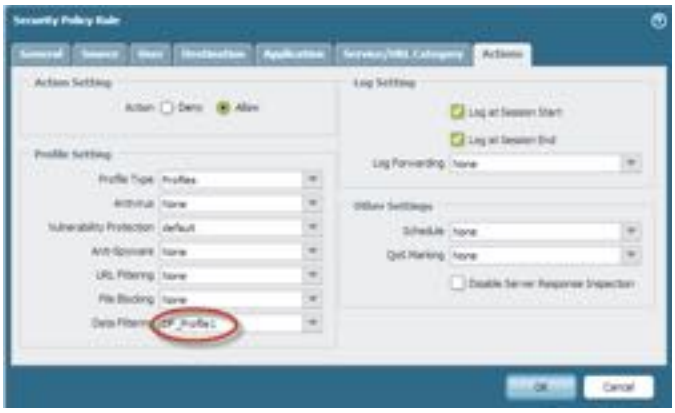 |
| Step 7 | **Commit** the configuration. | |
| Step 8 | Test the data filtering configuration.<br><br>If you have problems getting Data Filtering to work, you can check the Data Filtering log or the Traffic log to verify the application that you are testing with and make sure your test document has the appropriate number of unique Social Security Number instances. For example, an application such as Microsoft Outlook Web App may seem to be identified as web-browsing, but if you look at the logs, the application is **outlook-web**. Also increase the number of SSNs, or your custom pattern to make sure you are hitting the thresholds. | When testing, you must use real Social Security Numbers and each number must be unique. Also, when defining Custom Patterns as we did in this example with the word **confidential**, the pattern is case sensitive. To keep your test simple, you may want to just test using a data pattern first, then test the SSNs.<br>1. Access a client PC in the trust zone of the firewall and send an HTTP request to upload a .doc or .docx file that contains the exact information you defined for filtering.<br>2. Create a Microsoft Word document with one instance of the term confidential and five Social Security numbers with dashes.<br>3. Upload the file to a website. Use an HTTP site unless you have decryption configured, in which case you can use HTTPS.<br>4. Select **Monitoring > Logs > Data Filtering** logs.<br>5. Locate the log that corresponds to the file you just uploaded. To help filter the logs, use the source of your client PC and the destination of the web server. The action column in the log will show **reset-both**. You can now increase the number of Social Security Numbers in the document to test the block threshold. |

 © Palo Alto Networks, Inc.

## Set Up File Blocking

This example will describe the basic steps needed to set up file blocking and forwarding. In this configuration, we will configure the options needed to prompt users to continue before downloading .exe files from websites. When testing this example, be aware that you may have other systems between you and the source that may be blocking content.

| **Configure File Blocking** | | |
|---|---|---|
| Step 1 | Create the file blocking profile. | 1. Select **Objects > Security Profiles > File Blocking** and click **Add**.<br>2. Enter a **Name** for the file blocking profile, for example *Block_EXE*. Optionally enter a **Description**, such as *Block users from downloading exe files from websites.* |
| Step 2 | Configure the file blocking options. | 1. Click **Add** to define the profile settings.<br>2. Enter a **Name**, such as *BlockEXE*.<br>3. Set the **Applications** for filtering, for example **web-browsing**.<br>4. Set **File Types** to **exe**.<br>5. Set the **Direction** to **download**.<br>6. Set the **Action** to **continue**. By choosing the continue option, users will be prompted with a response page prompting them to click continue before the file will be downloaded.<br> <br>7. Click **OK** to save the profile. |
| Step 3 | Apply the file blocking profile to a security policy. | 1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Set Up Basic Security Policies.<br>2. Click the **Actions** tab within the policy rule.<br>3. In the Profile Settings section, click the drop-down and select the file blocking profile you configured. In this case, the profile name is *Block_EXE*.<br>4. **Commit** the configuration.<br>If no security profiles have been previously defined, select the **Profile Type** drop-down and select **Profiles**. You will then see the list of options to select the security profiles. |

**Configure File Blocking (Continued)**

Step 4 To test your file blocking configuration, access a client PC in the trust zone of the firewall and attempt to download a .exe file from a website in the untrust zone. A response page should display. Click **Continue** to download the file. You can also set other actions, such as alert only, forward (which will forward to WildFire), or block, which will not provide a continue page to the user. The following shows the default response page for File Blocking:

**Example: Default File Blocking Response Page**

Step 5 (Optional) Define custom file blocking response pages (**Device > Response Pages**). This allows you to provide more information to users when they see a response page. You can include information such as company policy information and contact information for a Helpdesk.

When you create a file blocking profile with the action continue or continue-and-forward (used for WildFire forwarding), you can only choose the application web-browsing. If you choose any other application, traffic that matches the security policy will not flow through the firewall due to the fact that the users will not be prompted with a continue page. Also, if the website uses HTTPS, you will need to have a decryption policy in place.

You may want to check your logs to confirm what application is being used when testing this feature. For example, if you are using Microsoft Sharepoint to download files, even though you are using a web-browser to access the site, the application is actually `sharepoint-base`, or `sharepoint-document`. You may want to set the application type to **Any** for testing.

 © Palo Alto Networks, Inc.

# Prevent Brute Force Attacks

A brute force attack uses a large volume of requests/responses from the same source or destination IP address to break into a system. The attacker employs a trial-and-error method to guess the response to a challenge or a request.

The Vulnerability Protection profile on the firewall includes signatures to protect you from brute force attacks. Each signature has an ID, Threat Name, Severity and is triggered when a pattern is recorded. The pattern specifies the conditions and interval at which the traffic is identified as a brute-force attack; some signatures are associated with another child signature that is of a lower severity and specifies the pattern to match against. When a pattern matches against the signature or child signature, it triggers the default action for the signature.

To enforce protection:

- Attach the vulnerability profile to a security rule. See Set Up Antivirus, Anti-spyware, and Vulnerability Protection.
- Install content updates that include new signatures to protect against emerging threats. See Manage Content Updates.

© Palo Alto Networks, Inc.

# Customize the Action and Trigger Conditions for a Brute Force Signature

The firewall includes two types of predefined brute force signatures—parent signature and child signature. A child signature is a single occurrence of a traffic pattern that matches the signature. A parent signature is associated with a child signature and is triggered when multiple events occur within a time interval and match the traffic pattern defined in the child signature.

Typically, a child signature is of default action *allow* because a single event is not indicative of an attack. In most cases, the action for a child signature is set to allow so that legitimate traffic is not blocked and threat logs are not generated for non-noteworthy events. Therefore, Palo Alto Networks recommends that you only change the default action after careful consideration.

In most cases, the brute force signature is a noteworthy event because of its recurrent pattern. If you would like to customize the action for a brute-force signature, you can do one of the following:

- Create a rule to modify the default action for all signatures in the brute force category. You can define the action to allow, alert, block, reset, or drop the traffic.
- Define an exception for a specific signature. For example, you can search for a CVE and define an exception for it.

  For a parent signature, you can modify both the trigger conditions and the action; for a child signature only the action can be modified.

To effectively mitigate an attack, the **block-ip address** action is recommended over the drop or reset action for most brute force signatures.

| **Customize the Threshold and Action for a Signature** | | |
|---|---|---|
| Step 1 | Create a new Vulnerability Protection Profile. | 1. Select **Objects > Security Profiles > Vulnerability Protection**.<br>2. Click **Add** and enter a **Name** for the Vulnerability Protection Profile. |
| Step 2 | Create a rule that defines the action for all signatures in a category.<br> | 1. Select **Rules**, click **Add** and enter a **Name** for the rule.<br>2. Set the **Action**. In this example, it is set to **Block**.<br>3. Set **Category** to **brute-force**.<br>4. (Optional) If blocking, specify whether to block server or client, the default is any.<br>5. See Step 3 to customize the action for a specific signature.<br>6. See Step 4 to customize the trigger threshold for a parent signature.<br>7. Click **OK** to save the rule and the profile. |

 © Palo Alto Networks, Inc.

**Customize the Threshold and Action for a Signature**

Step 3 (Optional) Customize the action for a specific signature.

1. Select **Exceptions** and click **Show all signatures** to find the signature you want to modify.

   To view all the signatures in the brute-force category, search for ( category contains 'brute-force' ).
2. To edit a specific signature, click the predefined default action in the **Action** column.

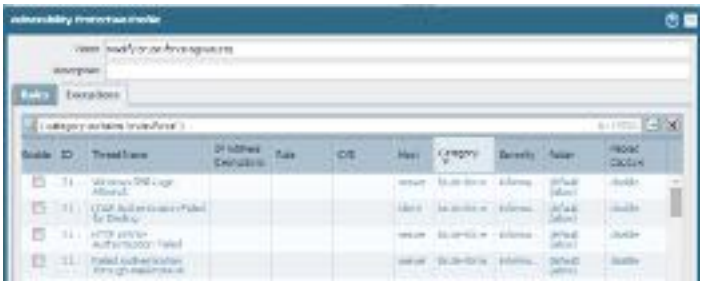

3. Set the action to **allow**, **alert** or **block-ip**.

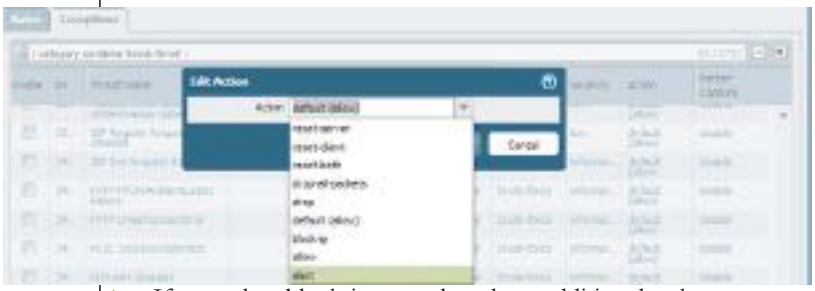

4. If you select block-ip, complete these additional tasks:
   a. Specify the **Time** period (in seconds) after which to trigger the action.
   b. In the **Track By** field, define whether to block the IP address by **IP source** or by **IP source and destination**.

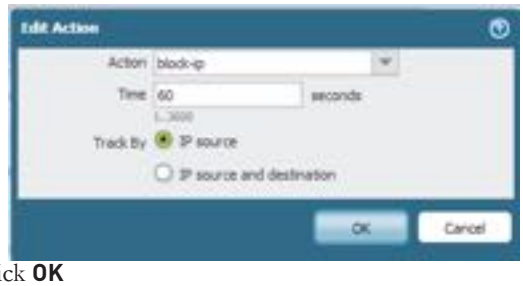


5. Click **OK**
6. For each modified signature, select the check box in the **Enable** column.
7. Click **OK**.

| Customize the Threshold and Action for a Signature | | |
|---|---|---|
| Step 4 | Customize the trigger conditions for a parent signature.<br><br>A parent signature that can be edited is marked with this icon: .<br><br>In this example, the search criteria was brute force category and CVE-2008-1447. | 1. Click to edit the time attribute and the aggregation criteria for the signature.<br> <br>2. To modify the trigger threshold specify the **Number of Hits** *per x* **seconds**.<br>3. Specify whether to aggregate the number of hits by **source**, **destination** or by **source and destination**.<br>4. Click **OK**. |
| Step 5 | Attach this new profile to a security rule. | |
| Step 6 | Save your changes. | 1. Click **Commit**. |

 © Palo Alto Networks, Inc.

# Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions

To monitor and protect your network from most Layer 4 and Layer 7 attacks, here are a few recommendations.

- Upgrade to the most current PAN-OS software version and content release version to ensure that you have the latest security updates. See Manage Content Updates and Install Software Updates.
- For servers, create security policy rules to only allow the application(s) that you sanction on each server. Verify that the standard port for the application matches the listening port on the server. For example, to ensure that only SMTP traffic is allowed to your email server set the Application to smtp and the Service to application-default. If your server uses only a subset of the standard ports (for example, if your SMTP server uses only port 587 while the smtp application has standard ports defined as 25 and 587), you should create a new custom service that only includes port 587 and use that new service in your security policy rule instead of using application-default. Additionally, make sure to restrict access to specific source and destinations zones and sets of IP addresses.
- Attach the following security profiles to your security policies to provide signature-based protection.
  - Create a vulnerability protection profile to block all vulnerabilities with severity low and higher.
  - Create an anti-spyware profile to block all spyware.
  - Create an antivirus profile to block all content that matches an antivirus signature.
- Block all unknown applications/traffic using security policy. Typically, the only applications that are classified as unknown traffic are internal or custom applications on your network, or potential threats. Because unknown traffic can be a non-compliant application or protocol that is anomalous or abnormal, or a known application that is using non-standard ports, unknown traffic should be blocked. See Manage Custom or Unknown Applications.
- Create a file blocking profile that blocks Portable Executable (PE) file types for Internet-based SMB (Server Message Block) traffic from traversing the trust to untrust zones, (ms-ds-smb applications).
- Create a zone protection profile that is configured to protect against packet-based attacks:
  - Remove TCP timestamps on SYN packets before the firewall forwards the packet—When you remove the TCP timestamp option in a SYN packet, the TCP stack on both ends of the TCP connection will not support TCP timestamps. Therefore, by disabling the TCP timestamp for a SYN packet, you can prevent an attack that uses different timestamps on multiple packets for the same sequence number.
  - Drop malformed packets.
  - Drop mismatched and overlapping TCP segments—By deliberately constructing connections with overlapping but different data in them, attackers can attempt to cause misinterpretation of the intent of the connection. This can be used to deliberately induce false positives or false negatives. An attacker can use IP spoofing and sequence number prediction to intercept a user's connection and inject his/her own data into the connection. PAN-OS uses this field to discard such frames with mismatched and overlapping data. The scenarios where the received segment will be discarded are when the segment received is contained within another segment, the segment received overlaps with part of another segment, or the segment completely contains another segment.

© Palo Alto Networks, Inc.

- Verify that support for IPv6 is enabled, if you have configured IPv6 addresses on your network hosts. (**Network > Interfaces > Ethernet> IPv6**)

  This allows access to IPv6 hosts and filters IPv6 packets that are encapsulated in IPv4 packets. Enabling support for IPv6 prevents IPv6 over IPv4 multicast addresses from being leveraged for network reconnaissance.

- Enable support for multicast traffic so that the firewall can enforce policy on multicast traffic. (**Network > Virtual Router > Multicast)**

- Enable the following CLI command to clear the URG bit flag in the TCP header and disallow out-of-band processing of packets.

  The urgent pointer in the TCP header is used to promote a packet for immediate processing by removing it from the processing queue and expediting it through the TCP/IP stack on the host. This process is called out-of-band processing. Because the implementation of the urgent pointer varies by host, to eliminate ambiguity, use the following CLI command to disallow out-of-band processing; the out-of-band byte in the payload becomes part of the payload and the packet is not processed urgently. Making this change allows you to remove ambiguity in how the packet is processed on the firewall and the host, and the firewall sees the exact same stream in the protocol stack as the host for whom the packet is destined.

  ```
  set deviceconfig setting tcp urgent-data clear
  ```

- If you configure the firewall to clear the URG bit flag and the packet has no other flags set in the TCP header, use the following CLI command to configure the firewall to drop packets with no flags:

  ```
  set deviceconfig setting tcp drop-zero-flag yes
  ```

- Enable the following CLI command for disabling the bypass-exceed-queue.

  The bypass exceed queue is required for out of order packets. This scenario is most common in an asymmetric environment where the firewall receives packets out of order. For identification of certain applications (App-ID) the firewall performs heuristic analysis. If the packets are received out of order, the data must be copied to a queue in order to complete the analysis for the application.

  ```
  set deviceconfig setting application bypass-exceed-queue no
  ```

- Enable the following CLI commands for disabling the inspection of packets when the out-of-order packet limit is reached. The Palo Alto Networks firewall can collect up to 64 out-of-order packets per session. This counter identifies that packets have exceeded the 64-packet limit. When the bypass setting is set to **no**, the device drops the out-of-order packets that exceed the 64-packet limit. A commit is required.

  ```
  set deviceconfig setting tcp bypass-exceed-oo-queue no
  set deviceconfig setting ctd tcp-bypass-exceed-queue no
  set deviceconfig setting ctd udp-bypass-exceed-queue no
  ```

- Enable the following CLI commands for checking the TCP timestamp. The TCP timestamp records when the segment was sent and allows the firewall to verify that the timestamp is valid for that session. Packets with invalid timestamps are dropped with this setting is enabled.

  ```
  set deviceconfig setting tcp check-timestamp-option yes
  ```

 © Palo Alto Networks, Inc.

- Disable the HTTP Range option. The HTTP Range option allows a client to fetch part of a file only. When a next-generation firewall in the path of a transfer identifies and drops a malicious file, it terminates the TCP session with a RST packet. If the web browser implements the HTTP Range option, it can start a new session to fetch only the remaining part of the file. This prevents the firewall from triggering the same signature again due to the lack of context into the initial session, while at the same time allowing the web browser to reassemble the file and deliver the malicious content. To prevent this, disable the HTTP Range option as follows:

```
set deviceconfig setting ctd skip-block-http-range no
```

# Enable Passive DNS Collection for Improved Threat Intelligence

Passive DNS is an opt-in feature that enables the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis in order to improve threat intelligence and threat prevention capabilities. The data collected includes non-recursive (i.e. originating from the local recursive resolver, not individual clients) DNS query and response packet payloads. Data submitted via the Passive DNS Monitoring feature consists solely of mappings of domain names to IP addresses. Palo Alto Networks retains no record of the source of this data and does not have the ability to associate it with the submitter at a future date.

The Palo Alto Networks threat research team uses this information to gain insight into malware propagation and evasion techniques that abuse the DNS system. Information gathered through this data collection is used to improve accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control signatures, and WildFire.

DNS responses are only forwarded to the Palo Alto Networks and will only occur when the following requirements are met:

- DNS response bit is set
- DNS truncated bit is not set
- DNS recursive bit is not set
- DNS response code is 0 or 3 (NX)
- DNS question count bigger than 0
- DNS Answer RR count is bigger than 0 or if it is 0, the flags need to be 3 (NX)
- DNS query record type are A, NS, CNAME, AAAA, MX

Passive DNS monitoring is disabled by default, but it is recommended that you enable it to facilitate enhanced threat intelligence. Use the following procedure to enable Passive DNS:

| **Enable Passive DNS** |
|---|
| 1. Select **Objects > Security Profiles > Anti-Spyware**. |
| 2. Select an existing profile to modify it or configure a new profile.<br>The Anti-Spyware profile must be attached to a security policy that governs your DNS server's external DNS traffic. |
| 3. Select the **DNS Signatures** tab and click the **Enable Passive DNS Monitoring** check box. |
| 4. Click **OK** and then **Commit**. |

 © Palo Alto Networks, Inc.

# Use DNS Queries to Identify Infected Hosts on the Network

The DNS sinkhole action in Anti-Spyware profiles enables the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. This allows you to identify hosts on your network that have been infected with malware. The following topics describe the DNS sinkhole action and provide instructions for enabling it and monitoring logs to identify infected hosts.


- DNS Sinkholing
- Configure DNS Sinkholing
- Identify Infected Hosts


© Palo Alto Networks, Inc.

## DNS Sinkholing

DNS sinkholing helps you to identify infected hosts on the protected network using DNS traffic in situations where the firewall cannot see the infected client's DNS query (that is, the firewall cannot see the originator of the DNS query). In a typical deployment where the firewall is north of the local DNS server, the threat log will identify the local DNS resolver as the source of the traffic rather than the actual infected host. Sinkholing malware DNS queries solves this visibility problem by forging responses to the client host queries directed at malicious domains, so that clients attempting to connect to malicious domains (for command-and-control, for example) will instead attempt to connect to a sinkhole IP address you define as illustrated in Configure DNS Sinkholing. Infected hosts can then be easily identified in the traffic logs because any host that attempts to connect to the sinkhole IP address are most likely infected with malware.

**Figure: DNS Sinkholing Example**

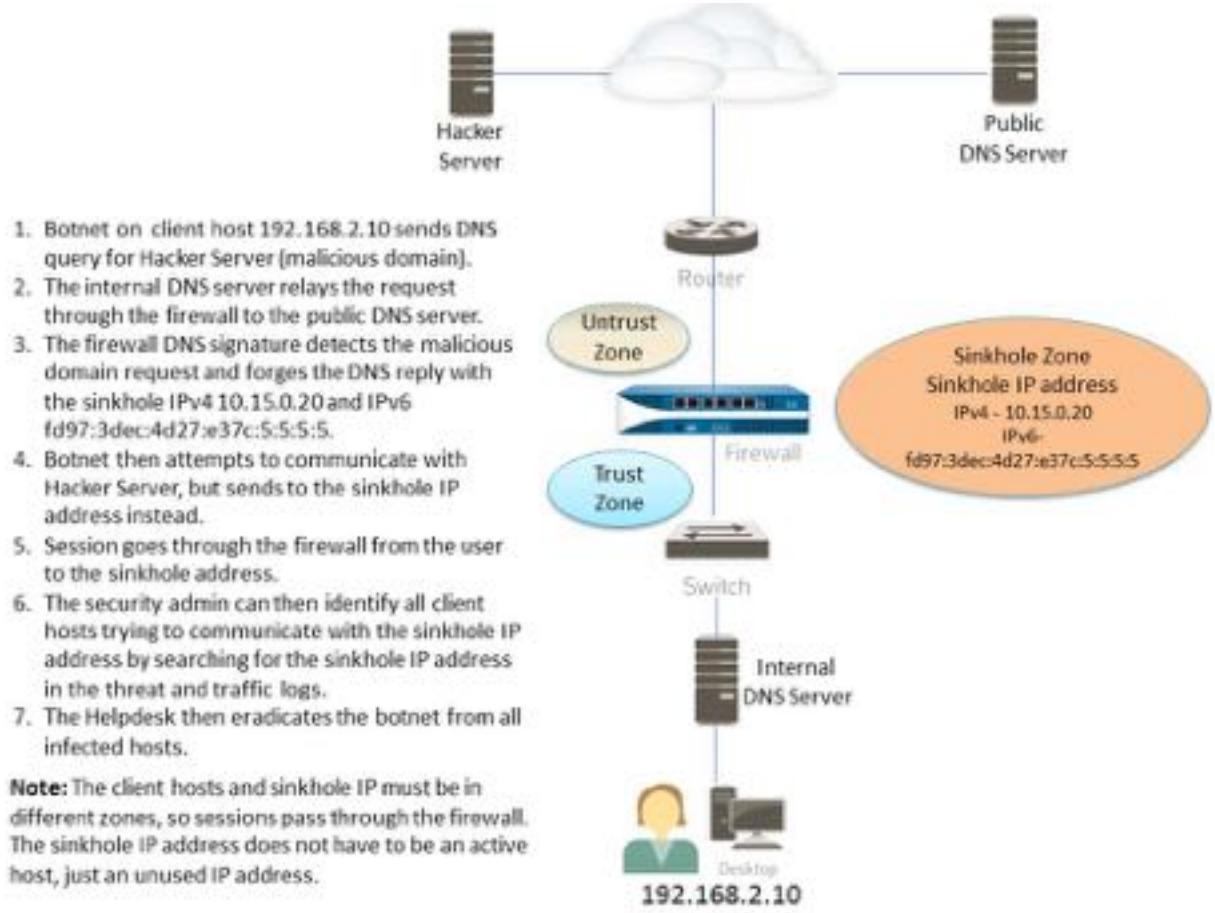


 © Palo Alto Networks, Inc.

## Configure DNS Sinkholing

To enable DNS Sinkholing, you must enable the action in an Anti-spyware profile and attach the profile to a security rule. When a client host attempts to access a malicious domain, the firewall forges the destination IP address in the packet using the IP address you configure as the DNS sinkhole address.

| Configure DNS Sinkholing | | |
|---|---|---|
| Step 1 | Obtain both an IPv4 and IPv6 address to use as the sinkhole IP addresses.<br><br>The DNS sinkhole address must be in a different zone than the client hosts to ensure that when an infected host attempts to start a session with the sinkhole IP address, it will be routed through the firewall. The reason both IPv4 and IPv6 are needed is because malicious software may perform DNS queries using one or both of these protocols.<br><br> This sinkhole addresses must be reserved for this purpose and do not need to be assigned to a physical host. You can optionally use a honey-pot server as a physical host to further analyze the malicious traffic. | This configuration example uses the following DNS sinkhole addresses:<br><br>IPv4 DNS sinkhole address—10.15.0.20<br><br>IPv6 DNS sinkhole address—fd97:3dec:4d27:e37c:5:5:5:5 |
| Step 2 | Configure the sinkhole interface and zone.<br><br>Traffic from the zone where the client hosts reside must route to the zone where the sinkhole IP address is defined, so traffic will be logged.<br><br> Use a dedicated zone for sinkhole traffic, because the infected host will be sending traffic to this zone. | 1. Select **Network > Interfaces** and select an interface to configure as your sinkhole interface.<br>2. In the **Interface Type** drop-down, select **Layer3**.<br>3. To add an IPv4 address, select the **IPv4** tab and select **Static** and then click **Add**. In this example, use the IPv4 address 10.15.0.20 as the sinkhole address.<br>4. Select the **IPv6** tab and click **Static** and then click **Add** and enter an IPv6 address and subnet mask. In this example, use the IPv6 address fd97:3dec:4d27:e37c::/64 as the sinkhole address.<br>5. Click **OK** to save.<br>6. To add a zone for the sinkhole, select **Network > Zones** and click **Add**.<br>7. Enter zone **Name**.<br>8. In the **Type** drop-down select **Layer3**.<br>9. In the **Interfaces** section, click **Add** and add the interface you just configured.<br>10. Click **OK**. |

| Configure DNS Sinkholing (Continued) | | |
|---|---|---|
| Step 3 | Enable DNS sinkholing on the anti-spyware profile. | 1. Select **Objects > Security Profiles > Anti-Spyware**.<br>2. Modify an existing profile, or select one of the existing defaults and clone it.<br>3. **Name** the profile and then select the **DNS Signatures** tab.<br>4. In the **Action on DNS queries** drop-down, select **sinkhole**.<br>5. In the **Sinkhole IPv4** field enter the sinkhole IPv4 sinkhole address you configured in Step 2 (10.15.0.20 in this example).<br>6. In the **Sinkhole IPv6** field enter the sinkhole IPv6 sinkhole address you configured in Step 2 (fd97:3dec:4d27:e37c:5:5:5:5 in this example).<br> The default sinkhole address is the loopback address (127.0.0.1 for IPv4 and ::1 for IPv6).<br>7. (Optional) In the **Packet Capture** drop-down, select **single-packet** or **extended-capture**. The single-packet option will capture the first packet of the session or you can select extended to capture between 1-50 packets. You can then use the packet captures for further analysis.<br>8. Click **OK** to save the profile. |
| Step 4 | Edit the security policy rule that allows traffic from client hosts in the trust zone to the untrust zone to include the sinkhole zone as a destination and attach the anti-spyware profile.<br><br>To ensure that you are identifying traffic from infected hosts, make these changes to the security rule(s) that allow traffic from client hosts in the trust zone to the untrust zone. By adding the sinkhole zone as a destination on the rule, you enable infected clients to send bogus DNS queries to the DNS sinkhole. | 1. Select **Policies > Security**.<br>2. Select an existing rule that allows traffic from the client host zone to the untrust zone.<br>3. On the **Destination** tab, **Add** the Sinkhole zone. This allows client host traffic to flow to the sinkhole zone.<br>4. On the **Actions** tab, select the **Log at Session Start** check box to enable logging. This will ensure that traffic from client hosts in the Trust zone will be logged when accessing the Untrust or Sinkhole zones.<br>5. In the **Profile Setting** section, select the **Anti-Spyware** profile in which you enabled DNS sinkholing.<br>6. Click **OK** to save the security rule and then **Commit**. |

 © Palo Alto Networks, Inc.

| Configure DNS Sinkholing (Continued) | | |
|---|---|---|
| Step 5 | To ensure that you will be able to identify infected hosts, verify that traffic going from the client host in the Trust zone to the new Sinkhole zone is being logged.<br><br>In this example, the infected client host is 192.168.2.10 and the Sinkhole IPv4 address is 10.15.0.20. | 1. From a client host in the trust zone, open a command prompt and run the following command:<br>`C:\>ping <sinkhole address>`<br>The following example output shows the ping request to the DNS sinkhole address at 10.15.0.2 and the result, which is `Request timed out` because in this example the sinkhole IP address is not assigned to a physical host:<br>`C:\>ping 10.15.0.20`<br>`Pinging 10.15.0.20 with 32 bytes of data:`<br>`Request timed out.`<br>`Request timed out.`<br>`Ping statistics for 10.15.0.20:`<br>`Packets: Sent = 4, Received = 0, Lost = 4 (100% loss)`<br>2. On the firewall, select **Monitor > Logs > Traffic** and find the log entry with the Source 192.168.2.10 and Destination 10.15.0.20. This will confirm that the traffic to the sinkhole IP address is traversing the firewall zones.<br> You can search and/or filter the logs and only show logs with the destination 10.15.0.20. To do this, click the IP address (10.15.0.20) in the **Destination** column, which will add the filter (addr.dst in 10.15.0.20) to the search field. Click the Apply Filter icon to the right of the search field to apply the filter. |
| Step 6 | Identify a malicious domain that you can use to verify that the DNS sinkhole functionality is configured properly.<br><br>You must test this feature using a malicious domain that is included in the firewall's current antivirus signature database. The DNS Signatures used to identify malicious domains is only part of the full antivirus signature database, which contains hundreds of thousands of signatures. | To find a malicious domain for testing:<br>1. Select **Device > Dynamic Updates** and in the **Antivirus** section click the **Release Notes** link for the current antivirus DB that is installed. You can also find the antivirus release notes on the support site in Dynamic Updates. In most cases, the signature update is an incremental update, so only new viruses and DNS signatures are listed. There are many antivirus signatures and DNS signatures that will already be installed on the firewall.<br>2. In the second column of the release note, locate a line item with a domain extension (for example, com, edu, or net). The left column will show the domain name. For example, in Antivirus release 1117-1560, there is an item in the left column named "tbsbana" and the right column lists "net".<br>The following shows the content in the release note for this line item:<br>`conficker:tbsbana1 variants: net`<br>Because this domain shows up in the current database, it will work for testing. |

| Configure DNS Sinkholing (Continued) | |
|---|---|
| Step 7 Test the sinkhole action<br><br>This is similar to the action that would be performed if the client host was infected and the malicious application was attempting to reach a hacker server using DNS queries. | 1. From the client host, open a command prompt.<br>2. Perform an NSLOOKUP to a URL that you identified as a known malicious domain in Step 6.<br>For example, using the URL `track.bidtrk.com`:<br>`C:\>nslookup track.bidtrk.com`<br>`Server: my-local-dns.local`<br>`Address: 10.0.0.222`<br>`Non-authoritative answer:`<br>`Name: track.bidtrk.com.org`<br>`Addresses: fd97:3dec:4d27:e37c:5:5:5:5`<br>`10.15.0.20`<br>In the output, note that the NSLOOKUP to the malicious domain has been forged using the sinkhole IP addresses that we configured (10.15.0.20). Because the domain matched a malicious DNS signature, the sinkhole action was performed.<br>3. Select **Monitor > Logs > Threat** and locate the corresponding threat log entry to verify that the correct action was taken on the NSLOOKUP request.<br>4. Perform a ping to `track.bidtrk.com`, which will generate network traffic to the sinkhole address. |

 © Palo Alto Networks, Inc.

## Identify Infected Hosts

After you have configured DNS sinkholing and verified that traffic to a malicious domain goes to the sinkhole address, you should regularly monitor traffic to the sinkhole address, so that you can track down the infected hosts and eliminate the threat.

| DNS Sinkhole Verification and Reporting | | |
|---|---|---|
| Step 1 | Use App Scope to identify infected client hosts. | 1. Select **Monitor > App Scope** and select **Threat Monitor**.<br>2. Click the **Show spyware** button along the top of the display page.<br>3. Select a time range.<br>The following screenshot shows three instances of Suspicious DNS queries, which were generated when the test client host performed an NSLOOKUP on a known malicious domain. Click the graph to see more details about the event.<br> |

**DNS Sinkhole Verification and Reporting (Continued)**

Step 2 Configure a custom report to identify all client hosts that have sent traffic to the sinkhole IP address, which is 10.15.0.20 in this example.

Forward to an SNMP manager, Syslog server and/or Panorama to enable alerts on these events.

In this example, the infected client host performed an NSLOOKUP to a known malicious domain that is listed in the Palo Alto Networks DNS Signature database. When this occurred, the query was sent to the local DNS server, which then forwarded the request through the firewall to an external DNS server. The firewall security policy with the Anti-Spyware profile configured matched the query to the DNS Signature database, which then forged the reply using the sinkhole address of 10.15.0.20 and fd97:3dec:4d27:e37c:5:5:5:5. The client attempts to start a session and the traffic log records the activity with the source host and the destination address, which is now directed to the forged sinkhole address.

Viewing the traffic log on the firewall allows you to identify any client host that is sending traffic to the sinkhole address. In this example, the logs show that the source address 192.168.2.10 sent the malicious DNS query. The host can then be found and cleaned. Without the DNS sinkhole option, the administrator would only see the local DNS server as the system that performed the query and would not see the client host that is infected. If you attempted to run a report on the threat log using the action "Sinkhole", the log would show the local DNS server, not the infected host.

1. Select **Monitor > Manage Custom Reports**.
2. Click **Add** and **Name** the report.
3. Define a custom report that captures traffic to the sinkhole address as follows:
   - **Database**—Select **Traffic Log**.
   - **Scheduled**—Enable **Scheduled** and the report will run every night.
   - **Time Frame**—30 days
   - **Selected Columns**—Select **Source address** or **Source User** (if you have User-ID configured), which will identify the infected client host in the report, and **Destination address**, which will be the sinkhole address.
   - In the section at the bottom of the screen, create a custom query for traffic to the sinkhole address (10.15.0.20 in this example). You can either enter the destination address in the **Query Builder** window (**addr.dst in 10.15.0.20**) or select the following in each column and click **Add**: Connector = and, Attribute = Destination Address, Operator = in, and Value = 10.15.0.20. Click **Add** to add the query.

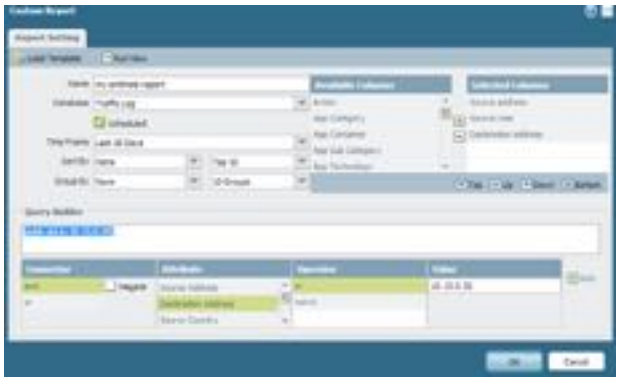

4. Click **Run Now** to run the report. The report will show all client hosts that have sent traffic to the sinkhole address, which indicates that they are most likely infected. You can now track down the hosts and check them for spyware.

5. To view scheduled reports that have run, select **Monitor > Reports**.

 © Palo Alto Networks, Inc.

# Content Delivery Network Infrastructure for Dynamic Updates

Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to the Palo Alto Networks devices. The devices access the web resources in the CDN to perform various App-ID and Content-ID functions. For enabling and scheduling the content updates, see Manage Content Updates.

The following table lists the web resources that the firewall accesses for a feature or application:

| Resource | URL | Static Addresses (If a static server is required) |
|---|---|---|
| Application Database | • updates.paloaltonetworks.com:443 | staticupdates.paloaltonetworks.com or the IP address 199.167.52.15 |
| Threat/Antivirus Database | • updates.paloaltonetworks.com:443<br>• downloads.paloaltonetworks.com:443<br>As a best practice, set the update server to updates.paloaltonetworks.com.This allows the Palo Alto Networks device to receive content updates from the server closest to it in the CDN infrastructure. | staticupdates.paloaltonetworks.com or the IP address 199.167.52.15 |
| PAN-DB URL Filtering | *.urlcloud.paloaltonetworks.com<br>Resolves to the primary URL s0000.urlcloud.paloaltonetworks.com and is then redirected to the regional server that is closest:<br>• s0100.urlcloud.paloaltonetworks.com<br>• s0200.urlcloud.paloaltonetworks.com<br>• s0300.urlcloud.paloaltonetworks.com<br>• s0500.urlcloud.paloaltonetworks.com | Static IP addresses are not available. However, you can manually resolve a URL to an IP address and allow access to the regional server IP address. |
| BrightCloud URL Filtering | • database.brightcloud.com:443/80<br>• service.brightcloud.com:80 | Contact BrightCloud Customer Support. |

| Resource | URL | Static Addresses (If a static server is required) |
|---|---|---|
| WildFire | • beta.wildfire.paloaltonetworks.com:443/80<br>• beta-s1.wildfire.paloaltonetworks.com:443/80<br>Beta sites are only accessed by a firewall running a Beta release version.<br>• mail.wildfire.paloaltonetworks.com:25<br>• wildfire.paloaltonetworks.com:443/80 | • mail.wildfire.paloaltonetworks.com:25 or the IP address 54.241.16.83<br>• wildfire.paloaltonetworks.com:443/80 or 54.241.8.199<br>• The regional URL/IP addresses are as follows:<br>• ca-s1.wildfire.paloaltonetworks.com:44 or 54.241.34.71<br>• va-s1.wildfire.paloaltonetworks.com:443or 174.129.24.252<br>• eu-s1.wildfire.paloaltonetworks.com:443 or 54.246.95.247<br>• sg-s1.wildfire.paloaltonetworks.com:443or 54.251.33.241<br>• jp-s1.wildfire.paloaltonetworks.com:443 or 54.238.53.161<br>• portal3.wildfire.paloaltonetworks.com:443/80 or 54.241.8.199<br>• ca-s3.wildfire.paloaltonetworks.com:443 or 54.241.34.71<br>• va-s3.wildfire.paloaltonetworks.com:443 or 23.21.208.35<br>• eu-s3.wildfire.paloaltonetworks.com:443 or 54.246.95.247<br>• sg-s3.wildfire.paloaltonetworks.com:443 or 54.251.33.241<br>• jp-s3.wildfire.paloaltonetworks.com:443 or 54.238.53.161<br>• wildfire.paloaltonetworks.com.jp:443/80 or 180.37.183.53<br>• wf1.wildfire.paloaltonetowrks.jp:443 or 180.37.180.37<br>• wf2.wildfire.paloaltonetworks.jp:443 or 180.37.181.18<br>• portal3.wildfire.paloaltonetworks.jp:443/80 or 180.37.183.53 |

 © Palo Alto Networks, Inc.

# Threat Prevention Resources

For more information on Threat Prevention, refer to the following sources:

- Creating Custom Threat Signatures
- Threat Prevention Deployment
- Understanding DoS Protection

To view a list of Threats and Applications that Palo Alto Networks products can identify, use the following links:

- Applipedia—Provides details on the applications that Palo Alto Networks can identify.
- Threat Vault—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

# Decryption

Palo Alto Networks firewalls provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Decryption on a Palo Alto Networks firewall includes the capability to enforce security policies on encrypted traffic, where otherwise the encrypted traffic might not be blocked and shaped according to your configured security settings. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted traffic. Enabling decryption on a Palo Alto Networks firewall can include preparing the keys and certificates required for decryption, creating a decryption policy, and configuring decryption port mirroring. See the following topics to learn about and configure decryption:

- Decryption Overview
- Decryption Concepts
- Configure SSL Forward Proxy
- Configure SSL Inbound Inspection
- Configure SSH Proxy
- Configure Decryption Exceptions
- Enable Users to Opt Out of SSL Decryption
- Configure Decryption Port Mirroring
- Temporarily Disable SSL Decryption

© Palo Alto Networks, Inc.

# Decryption Overview

Secure Sockets Layer (SSL) and Secure Shell (SSH) are encryption protocols used to secure traffic between two entities, such as a web server and a client. SSL and SSH encapsulate traffic, encrypting data so that it is meaningless to entities other than the client and server with the keys to decode the data and the certificates to affirm trust between the devices. Traffic that has been encrypted using the protocols SSL and SSH can be decrypted to ensure that these protocols are being used for the intended purposes only, and not to conceal unwanted activity or malicious content.

Palo Alto Networks firewalls decrypt encrypted traffic by using keys to transform strings (passwords and shared secrets) from ciphertext to plaintext (decryption) and from plaintext back to ciphertext (re-encrypting traffic as it exits the device). Certificates are used to establish the firewall as a trusted third party and to create a secure connection. SSL decryption (both forward proxy and inbound inspection) requires certificates to establish trust between two entities in order to secure an SSL/TLS connection. Certificates can also be used when excluding servers from SSL decryption. You can integrate a hardware security module (HSM) with a firewall to enable enhanced security for the private keys used in SSL forward proxy and SSL inbound inspection decryption. To learn more about storing and generating keys using an HSM and integrating an HSM with your firewall, see Secure Keys with a Hardware Security Module. SSH decryption does not require certificates.

Palo Alto Networks firewall decryption is policy-based, and can be used to decrypt, inspect, and control both inbound and outbound SSL and SSH connections. Decryption policies allow you to specify traffic for decryption according to destination, source, or URL category and in order to block or restrict the specified traffic according to your security settings. The firewall uses certificates and keys to decrypt the traffic specified by the policy to plaintext, and then enforces App-ID and security settings on the plaintext traffic, including Decryption, Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles. After traffic is decrypted and inspected on the firewall, the plaintext traffic is re-encrypted as it exits the firewall to ensure privacy and security. Use policy-based decryption on the firewall to achieve outcomes such as the following:

- Prevent malware concealed as encrypted traffic from being introduced into an corporate network.
- Prevent sensitive corporate information from moving outside the corporate network.
- Ensure the appropriate applications are running on a secure network.
- Selectively decrypt traffic; for example, exclude traffic for financial or healthcare sites from decryption by configuring a decryption exception.

The three decryption policies offered on the firewall, SSL Forward Proxy, SSL Inbound Inspection, and SSH Proxy, all provide methods to specifically target and inspect SSL outbound traffic, SSL inbound traffic, and SSH traffic, respectively. The decryption policies provide the settings for you to specify what traffic to decrypt and decryption profiles can be selected when creating a policy, in order to apply more granular security settings to decrypted traffic, such as checks for server certificates, unsupported modes, and failures. This policy-based decryption on the firewall gives you visibility into and control of SSL and SSH encrypted traffic according to configurable parameters.

You can also choose to extend a decryption configuration on the firewall to include Decryption Port Mirroring, which allows for decrypted traffic to be forwarded as plaintext to a third party solution for additional analysis and archiving.

 © Palo Alto Networks, Inc.

# Decryption Concepts

To learn about keys and certificates for decryption, decryption policies, and decryption port mirroring, see the following topics:

- Keys and Certificates for Decryption Policies
- SSL Forward Proxy
- SSL Inbound Inspection
- SSH Proxy
- Decryption Exceptions
- Decryption Port Mirroring

## Keys and Certificates for Decryption Policies

Keys are strings of numbers that are typically generated using a mathematical operation involving random numbers and large primes. Keys are used to transform other strings—such as passwords and shared secrets—from plaintext to ciphertext (called *encryption*) and from ciphertext to plaintext (called *decryption*). Keys can be symmetric (the same key is used to encrypt and decrypt) or asymmetric (one key is used for encryption and a mathematically related key is used for decryption). Any system can generate a key.

X.509 certificates are used to establish trust between a client and a server in order to establish an SSL connection. A client attempting to authenticate a server (or a server authenticating a client) knows the structure of the X.509 certificate and therefore knows how to extract identifying information about the server from fields within the certificate, such as its FQDN or IP address (called a *common name* or *CN* within the certificate) or the name of the organization, department, or user to which the certificate was issued. All certificates must be issued by a certificate authority (CA). After the CA verifies a client or server, the CA issues the certificate and signs it using its private key.

With a decryption policy configured, an SSL/TLS session between the client and the server is established only if the firewall trusts the CA that signed the server's certificate. In order to establish trust, the firewall must have the server's root CA certificate in its certificate trust list (CTL) and use the public key contained in that root CA certificate to verify the signature. The firewall then presents a copy of the server certificate signed by the Forward Trust certificate for the client to authenticate. You can also configure the firewall to use an enterprise CA as a forward trust certificate for SSL Forward Proxy. If the firewall does not have the server's root CA certificate in its CTL, the firewall will present a copy of the server certificate signed by the Forward Untrust certificate to the client. The Forward Untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites hosted by a server with untrusted certificates.

For detailed information on certificates, see Certificate Management.

To control the trusted CAs that your device trusts, use the **Device > Certificate Management > Certificates > Default Trusted Certificate Authorities** tab on the firewall web interface.

Table: Palo Alto Networks Device Keys and Certificates describes the different keys and certificates used by Palo Alto Networks devices for decryption. As a best practice, use different keys and certificates for each usage.

**Table: Palo Alto Networks Device Keys and Certificates**

| Key/Certificate Usage | Description |
|---|---|
| Forward Trust | The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall trusts. To configure a Forward Trust certificate on the firewall, see Step 2 in the Configure SSL Forward Proxy task. By default, the firewall determines the key size to use for the client certificate based on the key size of the destination server. However, you can also set a specific key size for the firewall to use. See Configure the Key Size for SSL Forward Proxy Server Certificates. For added security, store the forward trust certificate on a Hardware Security Module (HSM), see Store Private Keys on an HSM. |
| Forward Untrust | The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall does not trust. To configure a Forward Untrust certificate on the firewall, see Step 3 in the Configure SSL Forward Proxy task. |

 © Palo Alto Networks, Inc.

| Key/Certificate Usage | Description |
|---|---|
| SSL Exclude Certificate | Certificates for servers that you want to exclude from SSL decryption. For example, if you have SSL decryption enabled, but have certain servers that you do not want included in SSL decryption, such as the web services for your HR systems, you would import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Exclude a Server From Decryption. |
| SSL Inbound Inspection | The certificate used to decrypt inbound SSL traffic for inspection and policy enforcement. For this application, you would import the server certificate for the servers for which you are performing SSL inbound inspection, or store them on an HSM (see Store Private Keys on an HSM). |

## SSL Forward Proxy

Use an SSL Forward Proxy decryption policy to decrypt and inspect SSL/TLS traffic from internal users to the web. SSL Forward Proxy decryption prevents malware concealed as SSL encrypted traffic from being introduced to your corporate network; for example, if an employee is using her Gmail account from her corporate office and opens an email attachment that contains a virus, SSL Forward Proxy decryption will prevent the virus from infecting the client system and entering the corporate network.

With SSL Forward Proxy decryption, the firewall resides between the internal client and outside server. The firewall uses Forward Trust or Forward Untrust certificates to establish itself as a trusted third party to the session between the client and the server (For details on certificates, see Keys and Certificates for Decryption Policies). When the client initiates an SSL session with the server, the firewall intercepts the client's SSL request and forwards the SSL request to the server. The server sends a certificate intended for the client that is intercepted by the firewall. If the server's certificate is signed by a CA that the firewall trusts, the firewall creates a copy of the server's certificate signed by the Forward Trust certificate and sends the certificate to the client to authenticate. If the server's certificate is signed by a CA that the firewall does not trust, the firewall creates a copy of the server's certificate and signs it with the Forward Untrust certificate and sends it to the client. In this case, the client sees a block page warning that the site they're attempting to connect to is not trusted and the client can choose to proceed or terminate the session. When the client authenticates the certificate, the SSL session is established with the firewall functioning as a trusted forward proxy to the site that the client is accessing.

As the firewall continues to receive SSL traffic from the server that is destined for the client, it decrypts the SSL traffic into clear text traffic and applies security policies to the traffic. The traffic is then re-encrypted on the firewall and the firewall forwards the encrypted traffic to the client.

Figure: SSL Forward Proxy shows this process in detail.

 © Palo Alto Networks, Inc.

**Figure: SSL Forward Proxy**

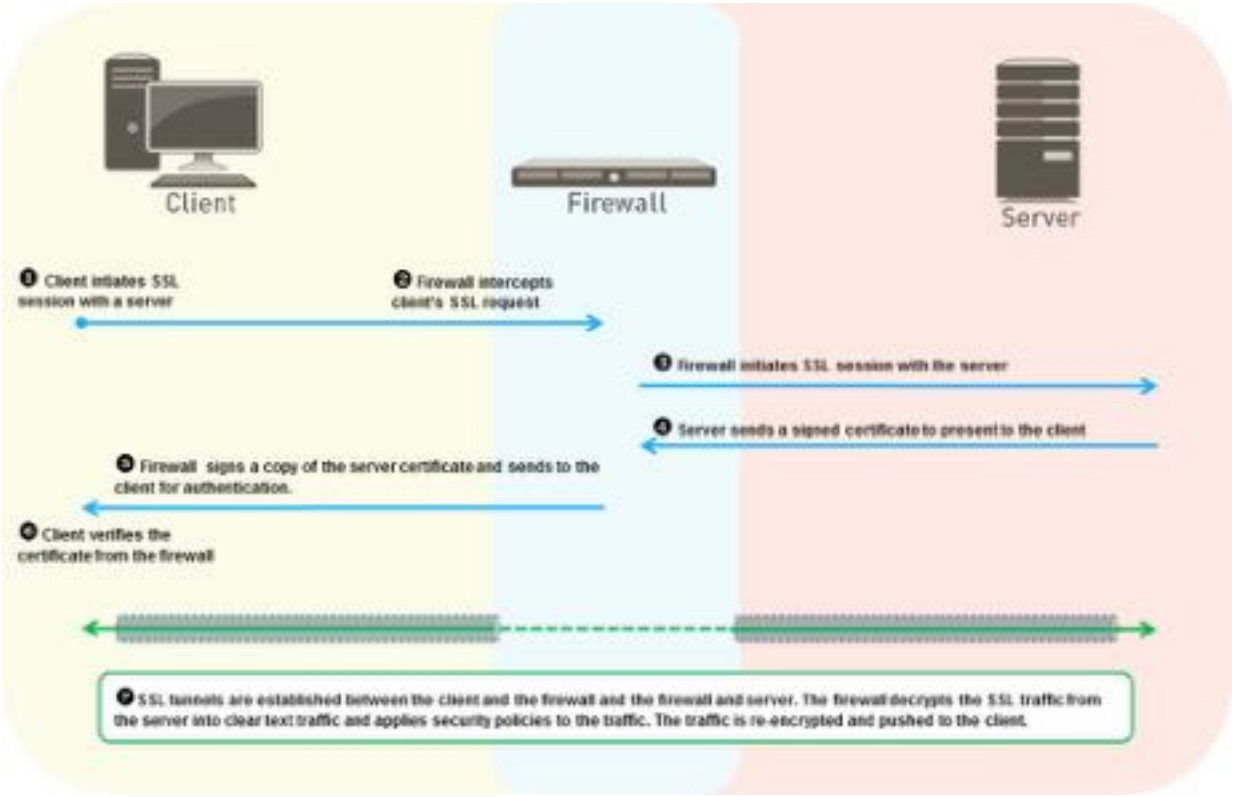


See Configure SSL Forward Proxy for details on configuring SSL Forward Proxy.

## SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL traffic from a client to a targeted server (any server you have the certificate for and can import it onto the firewall). For example, if an employee is remotely connected to a web server hosted on the company network and is attempting to add restricted internal documents to his Dropbox folder (which uses SSL for data transmission), SSL Inbound Inspection can be used to ensure that the sensitive data does not move outside the secure company network by blocking or restricting the session.

Configuring SSL Inbound Inspection includes importing the targeted server's certificate and key on to the firewall. Because the targeted server's certificate and key are imported on the firewall, the firewall is able to access the SSL session between the server and the client and decrypt and inspect traffic transparently, rather than functioning as a proxy. The firewall is able to apply security policies to the decrypted traffic, detecting malicious content and controlling applications running over this secure channel.

Figure: SSL Inbound Inspection shows this process in detail.

**Figure: SSL Inbound Inspection**

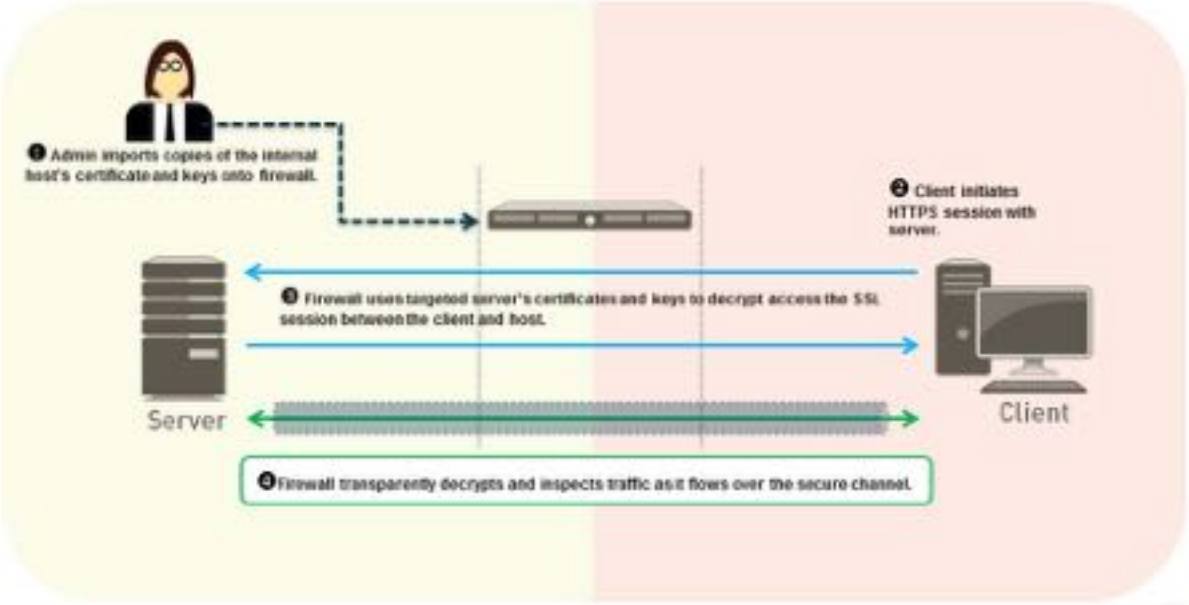


See Configure SSL Inbound Inspection for details on configuring SSL Inbound Inspection.

 © Palo Alto Networks, Inc.

## SSH Proxy

SSH Proxy provides the capability for the firewall to decrypt inbound and outbound SSH connections passing through the firewall, in order to ensure that SSH is not being used to tunnel unwanted applications and content. SSH decryption does not require any certificates and the key used for SSH decryption is automatically generated when the firewall boots up. During the boot up process, the firewall checks to see if there is an existing key. If not, a key is generated. This key is used for decrypting SSH sessions for all virtual systems configured on the device. The same key is also used for decrypting all SSH v2 sessions.

In an SSH Proxy configuration, the firewall resides between a client and a server. When the client sends an SSH request to the server, the firewall intercepts the request and forwards the SSH request to the server. The firewall then intercepts the server's response and forwards the response to the client, establishing an SSH tunnel between the firewall and the client and an SSH tunnel between the firewall and the server, with firewall functioning as a proxy. As traffic flows between the client and the server, the firewall is able to distinguish whether the SSH traffic is being routed normally or if it is using SSH tunneling (port forwarding). Content and threat inspections are not performed on SSH tunnels; however, if SSH tunnels are identified by the firewall, the SSH tunneled traffic is blocked and restricted according to configured security policies.

Figure: SSH Proxy Decryption shows this process in detail.

**Figure: SSH Proxy Decryption**

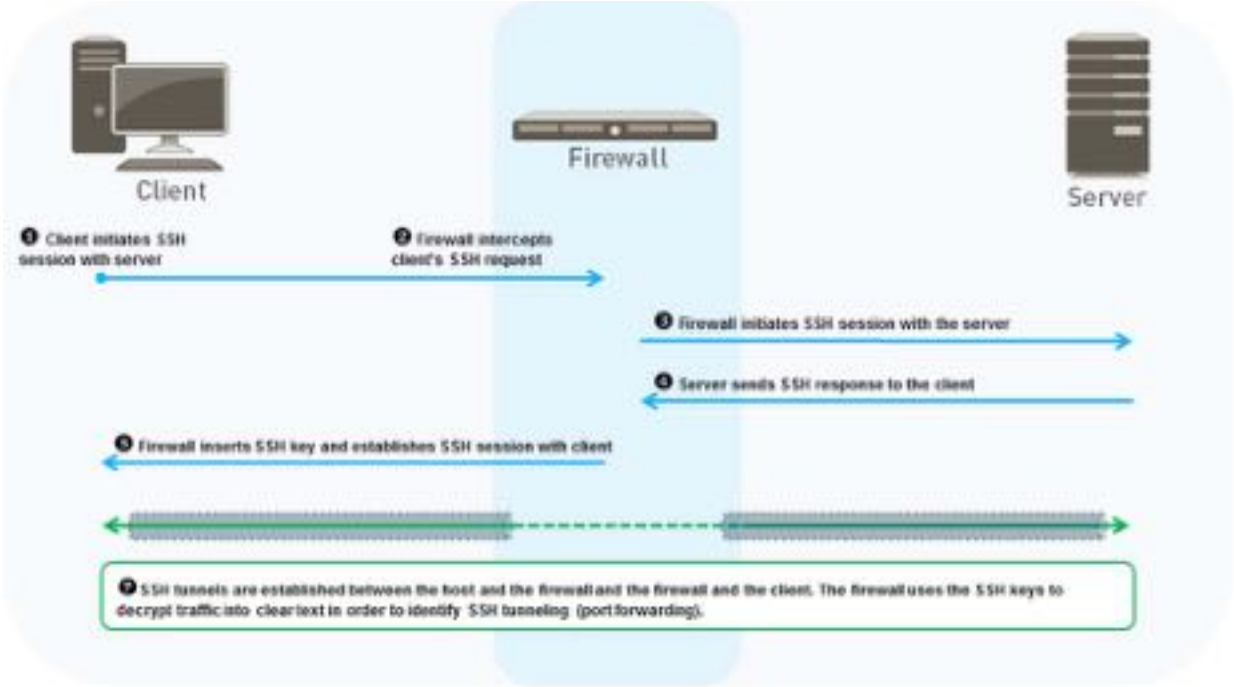


See Configure SSH Proxy for details on configuring an SSH Proxy policy.

## Decryption Exceptions

Traffic can also be excluded from decryption according to matching criteria (using a decryption policy), a targeted server's traffic can be excluded from decryption (using certificates), and some applications are excluded from decryption by default.

Applications that do not function properly when decrypted by the firewall and are automatically excluded from SSL decryption. The applications that are excluded from SSL decryption by default are excluded because these applications often fail when decrypted due to the application looking for specific details in the certificate that might not be present in the certificate generated for SSL Forward Proxy. Refer to the KB article List of Applications Excluded from SSL Decryption for a current list of applications excluded by default from SSL decryption on the firewall.

You can configure decryption exceptions for certain URL categories or applications that either do not work properly with decryption enabled or for any other reason, including for legal or privacy purposes. You can use a decryption policy to exclude traffic from decryption based on source, destination, service, and URL category. For example, with SSL decryption enabled, you can exclude traffic that is categorized as financial or health-related from decryption, using the URL category selection.

You can also exclude servers from SSL decryption based on the Common Name (CN) in the server's certificate. For example, if you have SSL decryption enabled but have certain servers that you do not want included in SSL decryption, such as the web services for your HR systems, you can exclude those servers from decryption by importing the server certificate onto the firewall and modifying the certificate to be an **SSL Exclude Certificate**.

To exclude traffic from decryption based on source, destination, service, URL category or to exclude a specific server's traffic from decryption, see Configure Decryption Exceptions.

 © Palo Alto Networks, Inc.

## Decryption Port Mirroring

The Decryption Port mirror feature provides the capability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures–such as NetWitness or Solera–for archiving and analysis. This feature is necessary for organizations that require comprehensive data capture for forensic and historical purposes or data leak prevention (DLP) functionality. Decryption port mirroring is available on PA-7050, PA-5000 Series and PA-3000 Series platforms only and requires that a free license be installed to enable this feature.

Keep in mind that the decryption, storage, inspection, and/or use of SSL traffic is governed in certain countries and user consent might be required in order to use the decryption port mirror feature. Additionally, use of this feature could enable malicious users with administrative access to the firewall to harvest usernames, passwords, social security numbers, credit card numbers, or other sensitive information submitted using an encrypted channel. Palo Alto Networks recommends that you consult with your corporate council before activating and using this feature in a production environment.

Figure: Decryption Port Mirroring shows the process for decryption port mirroring and the section Configure Decryption Port Mirroring describes how to license and use this feature.

**Figure: Decryption Port Mirroring**

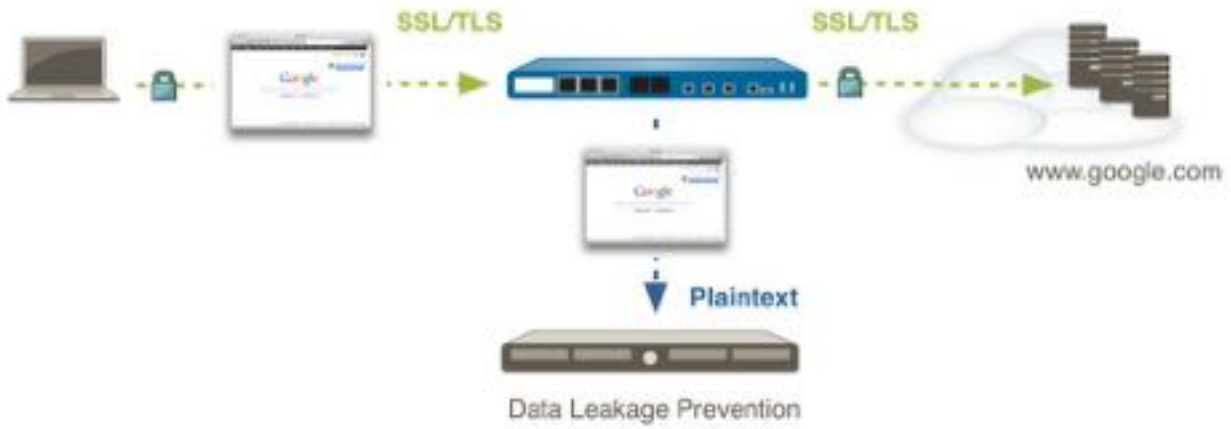

# Configure SSL Forward Proxy

Configuring SSL Forward Proxy decryption on the firewall requires setting up the certificates needed for SSL Forward Proxy decryption and creating an SSL Forward Proxy decryption policy. The firewall can use self-signed certificates or certificates signed by an enterprise CA to perform SSL Forward Proxy decryption.

By default, the firewall determines the key size to use for the client certificates it generates based on the key size of the destination server certificate. You can optionally set a static key size to use regardless of the key size of the destination server certificate. See Configure the Key Size for SSL Forward Proxy Server Certificates.

Use the following task to configure SSL Forward Proxy, including how to set up the certificates and create a decryption policy.

| **Configure SSL Forward Proxy** | | |
|---|---|---|
| Step 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. | View configured interfaces on the **Network > Interfaces > Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is. |


© Palo Alto Networks, Inc.

**Configure SSL Forward Proxy**

Step 2 Configure the forward trust certificate. Use either a self-signed certificate or a certificate signed by an enterprise CA.

**Using self-signed certificates**

When the certificate of the server that the client is connecting to is signed by a CA that is on the firewall's trusted CA list, the firewall signs a copy of the server's certificate with a self-signed forward trust certificate to present to the client for authentication. In this case, the self-signed certificate must be imported onto each client system so that the client recognizes the firewall as a trusted CA.

Use self-signed certificates for SSL Forward Proxy decryption if you do not use an enterprise CA or if you are only intended to perform decryption for a limited number of client systems (or if you are planning to use a centralized deployment).

**To use a self-signed certificate:**

1. Select **Device > Certificate Management > Certificates**.
2. Click **Generate** at the bottom of the window.
3. Enter a **Certificate Name**, such as *my-fwd-trust*.
4. Enter a **Common Name**, such as **192.168.2.1**. This should be the IP or FQDN that will appear in the certificate. In this case, we are using the IP of the trust interface. Avoid using spaces in this field.
5. Leave the **Signed By** field blank.
6. Click the **Certificate Authority** check box to enable the firewall to issue the certificate. Selecting this check box creates a certificate authority (CA) on the firewall that is imported to the client browsers, so clients trust the firewall as a CA.
7. Click **Generate** to generate the certificate.
8. Click the new certificate *my-fwd-trust* to modify it and enable the **Forward Trust Certificate** option.
9. Export the forward trust certificate for import into client systems by highlighting the certificate and clicking **Export** at the bottom of the window. Choose PEM format, and do not select the **Export private key** option. Because the certificate is self-signed, import it into the browser trusted root CA list on the client systems in order for the clients to trust it. When importing to the client browser, ensure the certificate is added to the Trusted Root Certification Authorities certificate store. On Windows systems, the default import location is the Personal certificate store. You can also simplify this process by using a centralized deployment, such as an Active Directory Group Policy Object (GPO).

   If the forward trust certificate is not imported on the client systems, users will see certificate warnings for each SSL site they visit.
10. Click **OK** to save.

## Configure SSL Forward Proxy

### Using an Enterprise CA

An enterprise CA can issue a signing certificate which the firewall can use to then sign the certificates for sites requiring SSL decryption. Send a Certificate Signing Request (CSR) for the enterprise CA to sign and validate. The firewall can then use the signed enterprise CA certificate for SSL Forward Proxy decryption. Because the enterprise CA is already trusted by the client systems, with this option, you do not need to distribute the certificate to client systems prior to configuring decryption.

**To use an enterprise CA signed certificate, generate a CSR:**

1. Select **Device > Certificate Management > Certificates** and click **Generate**.
2. Enter a **Certificate Name**, such as *my-fwd-proxy*.
3. In the **Signed By** drop-down, select **External Authority (CSR)**.
4. (Optional) If your enterprise CA requires it, add **Certificate Attributes** to further identify the firewall details, such as Country or Department.
5. Click **OK** to save the CSR. The pending certificate is now displayed on the **Device Certificates** tab.
6. Export the CSR:
   a. Select the pending certificate displayed on the **Device Certificates** tab.
   b. Click **Export** to download and save the certificate file.

   Leave **Export private key** unselected in order to ensure that the private key remains securely on the firewall.
   c. Click **OK**.
   d. Provide the certificate file to your enterprise CA. When you receive the signed enterprise CA certificate from your enterprise CA, save the signed enterprise CA certificate for import onto the firewall.
7. Import the signed enterprise CA onto the firewall:
   a. Select **Device > Certificate Management > Certificates** and click **Import**.
   b. Enter the pending **Certificate Name** exactly (in this case, *my-fwd-trust*). The **Certificate Name** that you enter must exactly match the pending certificate's name in order for the pending certificate to be validated.
   c. Select the signed **Certificate File** that you received from your enterprise CA.
   d. Click **OK**. The certificate is displayed as valid with the Key and CA check boxes selected.
   e. Select the validated certificate, in this case, *my-fwd-proxy*, to enable it as a **Forward Trust Certificate** to be used for SSL Forward Proxy decryption.
   f. Click **OK**.


© Palo Alto Networks, Inc.

| Configure SSL Forward Proxy | | |
|---|---|---|
| Step 3 | Configure the forward untrust certificate.<br><br>With SSL Forward Proxy decryption, when the site the client is connecting to uses a certificate signed by a CA that is not in the firewall's trusted CA list, the firewall presents a forward untrust certificate to the client. The forward untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites with untrusted certificates. | 1. Click **Generate** at the bottom of the certificates page.<br>2. Enter a **Certificate Name**, such as *my-fwd-untrust*.<br>3. Set the **Common Name**, for example **192.168.2.1**. Leave **Signed By** blank.<br>4. Click the **Certificate Authority** check box to enable the firewall to issue the certificate.<br>5. Click **Generate** to generate the certificate.<br>6. Click **OK** to save.<br>7. Click the new *my-ssl-fw-untrust* certificate to modify it and enable the **Forward Untrust Certificate** option.<br> Do not export the forward untrust certificate for import into client systems. If the forward trust certificate is imported on client systems, the users will not see certificate warnings for SSL sites with untrusted certificates.<br>8. Click **OK** to save. |
| Step 4 | (Optional) Set the key size of the SSL Forward Proxy certificates that the firewall presents to clients.<br><br> Changing the key size clears the current certificate cache. | 1. Select **Device > Setup > Session** and, in the Decryption Settings section, click **Forward Proxy Server Certificate Settings**.<br>2. Select a **Key Size**: **Defined by destination host** (default), **1024-bit RSA**, or **2048-bit RSA**. |
| Step 5 | (Optional) Create a Decryption profile.<br><br>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. An SSL Forward Proxy decryption profile can be used to perform checks for server certificates, unsupported modes, and failures and block or restrict traffic accordingly. For a complete list of checks that can be performed, navigate to **Objects > Decryption Profiles** on the firewall and click the help icon. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Select the **SSL Forward Proxy** tab to block and control specific aspects of SSL tunneled traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting **Block sessions if resources not available**.<br>3. Click **OK t**o save the profile. |

| Configure SSL Forward Proxy | |
|---|---|
| Step 6 Configure a decryption policy. | 1. Select **Policies > Decryption** and click **Add**.<br>2. On the **General** tab, give the policy a descriptive **Name**.<br>3. On the **Source** and **Destination** tabs, select **Any** for the **Source Zone** and **Destination Zone** to decrypt all SSL traffic destined for an external server. If you want to specify traffic from or to certain sources or destinations for decryption, click **Add**.<br>4. In the **URL Category** tab, leave **Any** to decrypt all traffic. If you only want to apply this profile to certain website categories, click **Add**.<br> Selecting a URL Category is useful when excluding certain sites from decryption. See Configure Decryption Exceptions.<br>5. On the **Options** tab, select **Decrypt** and select **SSL Forward Proxy** as the **Type** of decryption to perform.<br>6. (Optional) Select a **Decryption Profile** to apply additional settings to decrypted traffic (see Step 5).<br>7. Click **OK** to save. |
| Step 7 **Commit** the configuration. | With an SSL Forward Proxy decryption policy enabled, all traffic identified by the policy is decrypted. Decrypted traffic is blocked and restricted according to the profiles configured on the firewall (including the decryption profiles associated with the policy and Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles). Traffic is re-encrypted as it exits the firewall. |


© Palo Alto Networks, Inc.

# Configure SSL Inbound Inspection

Configuring SSL Inbound Inspection includes installing the targeted server's certificate on the firewall and creating an SSL Inbound Inspection decryption policy.

Use the following task to configure SSL Inbound Inspection.

| Configure SSL Inbound Inspection | | |
|---|---|---|
| Step 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. | View configured interfaces on the **Network > Interfaces > Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is. |
| Step 2 | Ensure that the targeted server's certificate is installed on the firewall. | On the web interface, select **Device > Certificate Management > Certificates > Device Certificates** to view certificates installed on the firewall.<br>To import the targeted server's certificate onto the firewall:<br>1. On the **Device Certificates** tab, select **Import**.<br>2. Enter a descriptive **Certificate Name**.<br>3. Browse for and select the targeted server's **Certificate File**.<br>4. Click **OK**. |
| Step 3 | (Optional) Create a Decryption profile.<br>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. An SSL Inbound Inspection decryption profile can be used to perform checks for unsupported modes and failures and block or restrict traffic accordingly. For a complete list of checks that can be performed, select **Objects > Decryption Profiles** and then click the help icon. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Select the **SSL Inbound Inspection** tab to block and control specific aspects of SSL traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting **Block sessions if resources not available**.<br>3. Click **OK t**o save the profile. |

| Configure SSL Inbound Inspection | |
|---|---|
| Step 4 Configure a decryption policy. | 1. Select **Policies > Decryption** and click **Add**.<br>2. On the **General** tab, give the policy a descriptive **Name**.<br>3. On the **Destination** tab, **Add** the **Destination Address** of the targeted server.<br>4. In the **URL Category** tab, leave **Any** to decrypt all traffic. If you only want to apply this profile to certain website categories, click **Add**.<br> Selecting a URL Category is useful when excluding certain sites from decryption. See Configure Decryption Exceptions.<br>5. On the **Options** tab, select **Decrypt** and select **SSL Inbound Inspection** as the **Type** of traffic to decrypt.<br>Select the **Certificate** for the internal server that is the destination of the inbound SSL traffic.<br>6. (Optional) Select a **Decryption Profile** to apply additional settings to decrypted traffic.<br>7. Click **OK** to save. |
| Step 5 **Commit** the configuration. | With an SSL Inbound Inspection decryption policy enabled, all SSL traffic identified by the policy is decrypted and inspected. Decrypted traffic is blocked and restricted according to the profiles configured on the firewall (including the decryption profiles associated with the policy and Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles). Traffic is re-encrypted as it exits the firewall. |


© Palo Alto Networks, Inc.

# Configure SSH Proxy

Configuring SSH Proxy does not require certificates and the key used to decrypt SSH sessions is generated automatically on the firewall during boot up.

Use the following task to configure SSH Proxy decryption.

| Configure SSH Proxy Decryption | | |
|---|---|---|
| Step 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. Decryption can only be performed on virtual wire, Layer 2, or Layer 3 interfaces. | View configured interfaces on the **Network > Interfaces > Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is. |
| Step 2 | (Optional) Create a Decryption profile.<br>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. The decryption profile can be used to perform checks for server certificates, unsupported modes, and failures and block or restrict traffic accordingly. For a complete list of checks that can be performed, navigate to **Objects > Decryption Profiles** on the firewall and then click the help icon. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Select the **SSH** tab to block and control specific aspects of SSH tunneled traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting **Block sessions if resources not available**.<br>3. Click **OK** to save the profile. |
| Step 3 | Configure a decryption policy. | 1. Select **Policies > Decryption** and click **Add**.<br>2. On the **General** tab, give the policy a descriptive **Name**.<br>3. On the **Source** and **Destination** tabs, select **Any** to decrypt all SSH traffic.<br>4. On the **URL Category** tab, select **Any** to decrypt all SSH traffic.<br>5. On the **Options** tab, select **Decrypt** and select **SSH Proxy** as the **Type** of traffic to decrypt.<br>6. (Optional) Select a **Decryption Profile** to apply additional settings to decrypted traffic.<br>7. Click **OK** to save. |
| Step 4 | **Commit** the configuration. | With the an SSH Proxy decryption policy enabled, all SSH traffic identified by the policy is decrypted and identified as either regular SSH traffic or as SSH tunneled traffic. SSH tunneled traffic is blocked and restricted according to the profiles configured on the firewall. Traffic is re-encrypted as it exits the firewall. |

# Configure Decryption Exceptions

You can purposefully exclude traffic from decryption based on matching criteria, such the traffic's source, destination, URL category, or service. You can also exclude a specific server's traffic from decryption. See the following topics to configure Decryption Exceptions:

- Exclude Traffic From Decryption
- Exclude a Server From Decryption

## Exclude Traffic From Decryption

To purposefully exclude applications or certain traffic from other existing SSL or SSH decryption policies, you can create a new decryption policy that defines the traffic to exclude from decryption with the **No Decrypt** action selected in the policy. You can define traffic for policy-based exclusion according to matching criteria, such as source, destination, URL categories, or the service (port or protocol). Make sure the decryption policy that excludes traffic from decryption is listed first in your decryption policy list by dragging and dropping the policy above the other decryption policies.

See the following procedure to configure a decryption policy that excludes traffic from SSL or SSH decryption.

| **Exclude Traffic from a Decryption Policy** | | |
|---|---|---|
| Step 1 | Create a decryption policy.<br><br>Use a decryption policy to exclude traffic from decryption according to the traffic's source and destination zones or addresses and URL categories. This example shows how to exclude traffic categorized as financial or health-related from SSL Forward Proxy decryption. | 1. Go to **Policies > Decryption** and click **Add**.<br>2. Give the policy a descriptive **Name**, such as *No-Decrypt-Finance-Health*.<br>3. On the **Source** and **Destination** tabs, select **Any** for the **Source Zone** and **Destination Zone** to apply the *No-Decrypt-Finance-Health* rule to all SSL traffic destined for an external server.<br>4. On the **URL Category** tab, **Add** the URL categories financial-services and health-and-medicine to the policy, specifying that traffic that matches these categories will not be decrypted.<br>5. On the **Options** tab, select **No Decrypt** and select the **Type** of decryption policy you are excluding the traffic from. For example, to exclude traffic categorized as financial or health-related from a separately configured SSL Forward Proxy decryption policy, select **SSL Forward Proxy** as the **Type**.<br>6. Click **OK** to save the *No-Decrypt-Finance-Health* decryption policy. |
| Step 2 | Move the decryption policy to the top of the list of decryption policies. | On the **Decryption > Policies** page, select the policy *No-Decrypt-Finance-Health*, and click **Move Up** until it appears at the top of the list (or you can drag and drop).<br><br>The order in which the decryption policies are listed is the order in which they are applied to network traffic. Moving the policy with the **No Decrypt** action applied to the top of the list ensures that the specified traffic is not decrypted according to another configured policy. |


© Palo Alto Networks, Inc.

| Exclude Traffic from a Decryption Policy | |
|---|---|
| Step 3 **Commit** the configuration. | A decryption policy with **No Decrypt** enabled ensures that the specified traffic is remains encrypted as it flows through the firewall, and that the traffic is not decrypted according to other decryption policies configured and listed on the **Policies > Decryption** page. |

## Exclude a Server From Decryption

You can exclude a targeted server's traffic from SSL decryption based on the Common Name (CN) in the server's certificate. For example, if you have SSL decryption enabled, you could configure a decryption exception for the server on your corporate network that hosts the web services for your HR systems. See the following procedure to configure a server's certificate so that the targeted server's traffic is excluded from decryption:

| Exclude a Server from Decryption | |
|---|---|
| Step 1 | Import the targeted server's certificate onto the firewall:<br>1. On the **Device > Certificate Management > Certificates > Device Certificates** tab, select **Import**.<br>2. Enter a descriptive **Certificate Name**.<br>3. Browse for and select the targeted server's **Certificate File**.<br>4. Click **OK**. |
| Step 2 | Select the targeted server's certificate on the **Device Certificates** tab and enable it as an SSL Exclude Certificate.<br>With the targeted server's certificate imported on the firewall and designated as an SSL Exclude Certificate, the server's traffic is not decrypted as it passes through the firewall. |

# Enable Users to Opt Out of SSL Decryption

In some cases, you might need to alert your users to the fact that the firewall is decrypting certain web traffic and allow them to terminate sessions that they do not want inspected. With SSL Opt Out enabled, the first time a user attempts to browse to an HTTPS site or application that matches your decryption policy, the firewall displays a response page notifying the user that it will decrypt the session. Users can either click **Yes** to allow decryption and continue to the site or click **No** to opt out of decryption and terminate the session. The choice to allow decryption applies to all HTTPS sites that users try to access for the next 24 hours, after which the firewall redisplays the response page. Users who opt out of SSL decryption cannot access the requested web page, or any other HTTPS site, for the next minute. After the minute elapses, the firewall redisplays the response page the next time the users attempt to access an HTTPS site.

The firewall includes a predefined SSL Decryption Opt-out Page that you can enable. You can optionally customize the page with your own text and/or images.

| **Enable Users to Opt Out of SSL Decryption** | |
|---|---|
| Step 1 (Optional) Customize the SSL Decryption Opt-out Page. | 1. Select **Device > Response Pages**.<br>2. Select the **SSL Decryption Opt-out Page** link.<br>3. Select the **Predefined** page and click **Export**.<br>4. Using the HTML text editor of your choice, edit the page.<br>5. If you want to add an image, host the image on a web server that is accessible from your end user systems.<br>6. Add a line to the HTML to point to the image. For example:<br>`<img src="http://cdn.slidesharecdn.com/Acme-logo-96x96.jpg?1382722588"/>`<br>7. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding.<br>8. Back on the firewall, select **Device > Response Pages**.<br>9. Select the **SSL Decryption Opt-out Page** link.<br>10. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.<br>11. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select shared to make it available to all virtual systems.<br>12. Click **OK** to import the file.<br>13. Select the response page you just imported and click **Close**. |
| Step 2 Enable SSL Decryption Opt Out. | 1. On the **Device > Response Pages** page, click the **Disabled** link.<br>2. Select the **Enable SSL Opt-out Page** and click **OK**.<br> <br>3. **Commit** the changes. |

 © Palo Alto Networks, Inc.

| Enable Users to Opt Out of SSL Decryption | |
|---|---|
| Step 3 Verify that the Opt Out page displays when you attempt to browse to a site. | From a browser, go to an encrypted site that matches your decryption policy.<br>Verify that the SSL Decryption Opt-out response page displays.<br> |

# Configure Decryption Port Mirroring

Before you can enable decryption port mirroring, you must obtain and install a Decryption Port Mirror license. The license is free of charge and can be activated through the support portal as described in the following procedure. After you install the Decryption Port Mirror license and reboot the firewall, you can enable decryption port mirroring.

Enabling decryption port mirroring includes enabling the forwarding of decrypted traffic and configuring a decrypt mirror interface. You can then create a decryption profile that specifies the interface and attach it to a decryption policy. To learn more about implementing Decryption Port Mirroring, see Decryption Port Mirroring.

Use the following procedure to obtain and install a Decryption Port Mirror license and configure Decryption Port Mirroring.

| **Configure Decryption Port Mirroring** | |
|---|---|
| Step 1 Request a license for each device on which you want to enable decryption port mirroring. | 1. Log in to the Palo Alto Networks Support site and navigate to the **Assets** tab.<br>2. Select the device entry for the device you want to license and select **Actions**.<br>3. Select **Decryption Port Mirror**. A legal notice displays.<br>4. If you are clear about the potential legal implications and requirements, click **I understand and wish to proceed**.<br>5. Click **Activate**.<br> 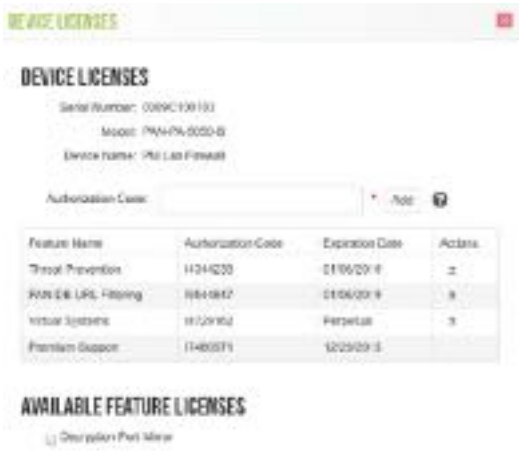 |

 © Palo Alto Networks, Inc.

| Configure Decryption Port Mirroring (Continued) | | |
|---|---|---|
| Step 2 | Install the Decryption Port Mirror license on firewall. | 1. From the firewall's web interface, select **Device > Licenses**.<br>2. Click **Retrieve license keys from license server**.<br>3. Verify that the license has been activated on the firewall.<br> <br><br>4. Reboot the firewall (**Device > Setup > Operations**). This feature will not be available for configuration until PAN-OS reloads. |
| Step 3 | Enable the ability to mirror decrypted traffic. Superuser permission is required to perform this step. | **On a firewall with a single virtual system:**<br>1. Select **Device > Setup > Content - ID**.<br>2. Select the **Allow forwarding of decrypted content** check box.<br>3. Click **OK** to save.<br>**On a firewall with multiple virtual systems:**<br>1. Select **Device > Virtual System.**<br>2. Select a Virtual System to edit or create a new Virtual System by selecting **Add**.<br>3. Select the **Allow forwarding of decrypted content** check box.<br>4. Click **OK** to save. |
| Step 4 | Configure a decrypt mirror interface. | 1. Select **Network > Interfaces > Ethernet**.<br>2. Select the Ethernet interface that you want to configure for decryption port mirroring.<br>3. Select **Decrypt Mirror** as the **Interface Type**.<br>This interface type will appear only if the Decryption Port Mirror license is installed.<br>4. Click **OK** to save. |

© Palo Alto Networks, Inc.

| Configure Decryption Port Mirroring (Continued) | | |
|---|---|---|
| Step 5 | Configure a Decryption Profile to enable decryption port mirroring. | 1. Select **Objects > Decryption Profile**.<br>2. Select the **Interface** to use for **Decryption Mirroring**.<br>The **Interface** drop-down contains all Ethernet interfaces that have been defined as the type: **Decrypt Mirror**.<br>3. Specify whether to mirror decrypted traffic before or after policy enforcement.<br>By default, the firewall will mirror all decrypted traffic to the interface before security policies lookup, which allows you to replay events and analyze traffic that generates a threat or triggers a drop action. If you want to only mirror decrypted traffic after security policy enforcement, select the **Forwarded Only** check box. With this option, only traffic that is forwarded through the firewall is mirrored. This option is useful if you are forwarding the decrypted traffic to other threat detection devices, such as a DLP device or another intrusion prevention system (IPS).<br>4. Click **OK** to save the decryption profile. |
| Step 6 | Set a decryption policy for decryption port mirroring. | 1. Select **Policies > Decryption**.<br>2. Click **Add** to configure a decryption policy or select an existing decryption policy to edit.<br>3. In the **Options** tab, select **Decrypt** and the **Decryption Profile** created in Step 4.<br>4. Click **OK** to save the policy. |
| Step 7 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# Temporarily Disable SSL Decryption

In some cases you may want to temporarily disable SSL decryption. For example, if your users are having problems accessing an encrypted site or application, you may want to disable SSL decryption in order to troubleshoot the issue. Although you could disable the associated decryption policies, modifying the policies is a configuration change that requires a Commit. Instead, use the following command to temporarily disable SSL decryption and then re-enable it after you finish troubleshooting. This command does not require a commit and it does not persist in your configuration after a reboot.

| **Temporarily Disable SSL Decryption** | |
|---|---|
| • Disable SSL Decryption | `set system setting ssl-decrypt skip-ssl-decrypt yes` |
| • Re-enable SSL Decryption | `set system setting ssl-decrypt skip-ssl-decrypt no` |

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

# URL Filtering

The Palo Alto Networks URL filtering solution is a powerful PAN-OS feature that is used to monitor and control how users access the web over HTTP and HTTPS. The following topics provide an overview of URL filtering, configuration and troubleshooting information, and best practices for getting the most out of this feature:


- URL Filtering Overview
- URL Filtering Concepts
- PAN-DB Categorization Workflow
- Enable a URL Filtering Vendor
- Determine URL Filtering Policy Requirements
- Monitor Web Activity
- Configure URL Filtering
- Log Client IP Addresses for Proxied HTTP/HTTPS Requests
- Customize the URL Filtering Response Pages
- Configure URL Admin Override
- Enable Safe Search Enforcement
- URL Filtering Use Case Examples
- Troubleshoot URL Filtering


© Palo Alto Networks, Inc.

# URL Filtering Overview

The Palo Alto Networks URL filtering feature compliments the App-ID feature by enabling you to configure your firewall to identify and control access to web (HTTP and HTTPS) traffic. By implementing URL filtering profiles in security policies and by using URL categories as a match criteria in policies (captive portal, decryption, security, and QoS), you will gain complete visibility and control of the traffic that traverses your firewall and will be able to safely enable and control how your users access the web.

The Palo Alto Networks URL filtering solution utilizes a URL filtering database that contains millions of websites and each website is placed in one of approximately 60 different categories. A URL filtering profile that contains the list of categories is then applied to a security policy that allows web traffic (HTTP/HTTPS) from the internal users to the Internet. After the URL filtering profile is applied and the alert or block action is set on a category, you will gain complete visibility into the websites that users access and can then decide which websites or website categories should be allowed, blocked, or logged. You can also define a list of URLs in the URL filtering profile that will always be blocked or allowed and you can create custom URL categories that contain a list of URLs that can be used the same way as the default category list. These same URL categories can also be used as a match criteria in other policies, such as captive portal, decryption, and QoS.

▲ URL Filtering Vendors

▲ Interaction Between App-ID and URL Categories

## URL Filtering Vendors

Palo Alto Networks firewalls support two vendors for URL filtering purposes:

- **PAN-DB**—A Palo Alto Networks developed URL filtering database that is tightly integrated into PAN-OS by utilizing high-performance local caching to perform maximum inline performance for URL lookups while a distributed cloud architecture provides coverage for the latest websites. In addition, PAN-DB is tightly integrated with WildFire such that whenever WildFire deems a site malicious, it updates the corresponding PAN-DB URL category to malware, immediately blocking any future access to the site as long as you have a URL Filtering profile attached to the security policy rule. To view a list of PAN-DB URL filtering categories, refer to https://urlfiltering.paloaltonetworks.com/CategoryList.aspx.
- **BrightCloud**—A third-party URL database that is owned by Webroot, Inc. and is integrated into PAN-OS firewalls. For information on the BrightCloud URL database, visit http://brightcloud.com.

For instructions on configuring the firewall to use one of the supported URL Filtering vendors, see Enable a URL Filtering Vendor.

 © Palo Alto Networks, Inc.

## Interaction Between App-ID and URL Categories

The Palo Alto Networks URL filtering feature in combination with Application Identification (App-ID) provides unprecedented protection against a full spectrum of legal, regulatory, productivity, and resource utilization risks. While App-ID gives you control over what applications users can access, URL filtering provides control over related web activity. When combined with User-ID, you can also apply these controls based on users and groups.

With today's application landscape and the way many applications use HTTP and HTTPS, you will need to determine when to use App-ID and when to use URL filtering in order to define comprehensive web access policies. In most cases, if an App-ID signature exists, you will want to use App-ID to control the content at the application level. For example, although you can control access to Facebook and/or LinkedIn using URL filtering, this would not block the use of all related applications, such as email, chat, as well as any new application that is introduced after you implement your policy.

In some cases, you will want to use both URL filtering and App-ID, but to ensure that conflicts do not occur in your policies, it is important to understand how these features work together. Palo Alto Networks generates signatures for many applications and those signatures can be very granular in regard to various features within the web-based applications—whereas URL filtering would only apply actions based on a specific website or URL category. For example, you may want to block social networking sites in general, but want to allow a few sites in that category to be accessible to specific departments and then control what features of the website are available to the allowed users. For more information, see URL Filtering Use Case Examples.

© Palo Alto Networks, Inc.

# URL Filtering Concepts


- URL Categories
- URL Filtering Profile
- URL Filtering Profile Actions
- Block and Allow Lists
- Safe Search Enforcement
- Container Pages
- HTTP Header Logging
- URL Filtering Response Pages
- URL Category as Policy Match Criteria



© Palo Alto Networks, Inc.

## URL Categories

Each website defined in the URL filtering database is assigned one of approximately 60 different URL categories. There are two ways to make use of URL categorization on the firewall:

- **Block or allow traffic based on URL category**—You can create a URL Filtering Profile that specifies an action for each URL category and attach the profile to a policy. Traffic that matches the policy would then be subject to the URL filtering settings in the profile. For example, to block all gaming websites you would set the block action for the URL category *games* in the URL profile and attach it to the security policy rule(s) that allow web access. See Configure URL Filtering for more information.
- **Match traffic based on URL category for policy enforcement**—If you want a specific policy rule to apply only to web traffic to sites in a specific category, you would add the category as match criteria when you create the policy rule. For example, you could use the URL category *streaming-media* in a QoS policy to apply bandwidth controls to all websites that are categorized as streaming media. See URL Category as Policy Match Criteria for more information.

By grouping websites into categories, it makes it easy to define actions based on certain types of websites. In addition to the standard URL categories, there are three additional categories:

| Categories | Description |
|---|---|
| Not-resolved | Indicates that the website was not found in the local URL filtering database and the firewall was unable to connect to the cloud database to check the category. When a URL category lookup is performed, the firewall first checks the dataplane cache for the URL; if no match is found, it will then check the management plane cache, and if no match is found there, it queries the URL database in the cloud.<br>When deciding on what action to take for traffic that is categorized as *not-resolved*, be aware that setting the action to block may be very disruptive to users.<br>For more information on troubleshooting lookup issues, see Troubleshoot URL Filtering. |
| Private-ip-addresses | Indicates that the website is a single domain (no sub-domains), the IP address is in the private IP range, or the URL root domain is unknown to the cloud. |
| Unknown | The website has not yet been categorized, so it does not exist in the URL filtering database on the firewall or in the URL cloud database.<br>When deciding on what action to take for traffic categorized as *unknown*, be aware that setting the action to block may be very disruptive to users because there could be a lot of valid sites that are not in the URL database yet. If you do want a very strict policy, you could block this category, so websites that do not exist in the URL database cannot be accessed.<br>See Configure URL Filtering. |

## URL Filtering Profile

A URL filtering profile is a collection of URL filtering controls that are applied to individual security policy rules to enforce your web access policy. The firewall comes with a default profile that is configured to block threat-prone categories, such as malware, phishing, and adult. You can use the default profile in a security policy, clone it to be used as a starting point for new URL filtering profiles, or add a new URL profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories. For example, you may want to block social-networking sites, but allow some websites that are part of the social-networking category.

The following topics describe the components of the URL Filtering profile:

- URL Filtering Profile Actions
- Block and Allow Lists
- Safe Search Enforcement
- Container Pages
- HTTP Header Logging

 © Palo Alto Networks, Inc.

## URL Filtering Profile Actions

The URL Filtering profile specifies an action for each known URL category. By default, when you create a new URL Filtering profile, the action is set to **allow** for all categories. This means that the users will be able to browse to all sites freely and will not be logged.

The default URL Filtering profile is set to allow access to all URL categories except for the following threat-prone categories, which are blocked: abused-drugs, adult, gambling, hacking, malware, phishing, questionable, and weapons. As a best practice, when you need a custom URL Filtering profile, clone the default profile rather than creating a new one to preserve these settings.

| Action | Description |
|---|---|
| **alert** | The website is allowed and a log entry is generated in the URL filtering log. |
| **allow** | The website is allowed and no log entry is generated. |
| **block** | The website is blocked and the user will see a response page and will not be able to continue to the website. A log entry is generated in the URL filtering log. |
| **continue** | The user will be prompted with a response page indicating that the site has been blocked due to company policy, but the user is prompted with the option to continue to the website. The **continue** action is typically used for categories that are considered benign and is used to improve the user experience by giving them the option to continue if they feel the site is incorrectly categorized. The response page message can be customized to contain details specific to your company. A log entry is generated in the URL filtering log.<br> The Continue page will not be displayed properly on client machines that are configured to use a proxy server. |
| **override** | The user will see a response page indicating that a password is required to allow access to websites in the given category. With this option, the security admin or helpdesk person would provide a password that will grant temporary access to all websites in the given category. A log entry is generated in the URL filtering log. See Configure URL Admin Override.<br> The Override page will not be displayed properly on client machines that are configured to use a proxy server. |
| **none** | The **none** action only applies to custom URL categories. Select **none** to ensure that if multiple URL profiles exist, the custom category will not have any impact on other profiles. For example, if you have two URL profiles and the custom URL category is set to **block** in one of the profiles, the other profile should have the action set to **none** if you do not want it to apply.<br>Also, in order to delete a custom URL category, it must be set to **none** in any profile where it is used. |

## Block and Allow Lists

Block and allow lists allow you to define specific URLs or IP addresses in the URL filtering profile that are always allowed or always blocked, regardless of the action defined for the URL category. When entering URLs in the *Block List* or *Allow List*, enter each URL or IP address in a new row separated by a new line. When using wildcards in the URLs, follow these rules:

- Do not include HTTP and HTTPS when defining URLs. For example, enter `www.paloaltonetworks.com` or `paloaltonetworks.com` instead of https://www.paloaltonetworks.com.
- Entries in the block list must be an exact match and are case-insensitive.

  For example: If you want to prevent a user from accessing any website within the domain paloaltonetworks.com, you would also enter `*.paloaltonetworks.com`, so whatever domain prefix (http://, www, or a sub-domain prefix such as mail.paloaltonetworks.com) is added to the address, the specified action will be taken. The same applies to the sub-domain suffix; if you want to block paloaltonetworks.com/en/US, you would need to add `paloaltonetworks.com/*` as well.

  Further, if you want to specify that the match is for `paloaltonetworks.com` and not `paloaltonetworks.com.au`, you must add a `/` after `.com` to specify the scope for the match to just be `.com` and not for the domain suffix `.com.au`. In this case, you need to add the entry as `*.paloaltonetworks.com/`

  Block and allow lists support wildcard patterns. Use the following characters as separators:

  `.`

  `/`

  `?`

  `&`

  `=`

  `;`

  `+`

  Every substring that is separated by the characters listed above is considered a token. A token can be any number of ASCII characters that does not contain any separator character or *. For example, the following patterns are valid:

  *.yahoo.com (tokens are: "*", "yahoo" and "com")

  www.*.com (tokens are: "www", "*" and "com")

  www.yahoo.com/search=* (tokens are: "www", "yahoo", "com", "search", "*")

  The following patterns are invalid because the character "*" is not the only character in the token.

  ww*.yahoo.com

  www.y*.com

 © Palo Alto Networks, Inc.

## Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos in search query return traffic. On the firewall, you can Enable Safe Search Enforcement so that the firewall will block search results if the end user is not using the strictest safe search settings in the search query. The firewall can enforce safe search for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. This is a best-effort setting and is not guaranteed by the search providers to work with every website.

To use this feature you must enable the **Safe Search Enforcement** option in a URL filtering profile and attach it to a security policy. The firewall will then block any matching search query return traffic that is not using the strictest safe search settings. There are two methods for blocking the search results:

- Block Search Results that are not Using Strict Safe Search Settings—When an end user attempts to perform a search without first enabling the strictest safe search settings, the firewall blocks the search query results and displays the URL Filtering Safe Search Block Page. By default, this page will provide a URL to the search provider settings for configuring safe search.
- Enable Transparent Safe Search Enforcement—When an end user attempts to perform a search without first enabling the strict safe search settings, the firewall blocks the search results with an HTTP 503 status code and redirects the search query to a URL that includes the safe search parameters. You enable this functionality by importing a new URL Filtering Safe Search Block Page containing the Javascript for rewriting the search URL to include the strict safe search parameters. In this configuration, users will not see the block page, but will instead be automatically redirected to a search query that enforces the strictest safe search options. This safe search enforcement method requires Content Release version 475 or later and is only supported for Google, Yahoo, and Bing searches.

Also, because most search providers now use SSL to return search results, you must also configure a Decryption policy for the search traffic to enable the firewall to inspect the search traffic and enforce safe search.

Safe search enforcement enhancements and support for new search providers is periodically added in content releases. This information is detailed in the Application and Threat Content Release Notes. How sites are judged to be safe or unsafe is performed by each search provider, not by Palo Alto Networks.

Safe search settings differ by search provider as detailed in Table: Search Provider Safe Search Settings.

**Table: Search Provider Safe Search Settings**

| Search Provider | Safe Search Setting Description |
|---|---|
| Google/YouTube | Offers safe search on individual computers or network-wide through Google's safe search virtual IP address:<br><br>**Safe Search Enforcement for Google Searches on Individual Computers**<br><br>In the Google Search Settings, the **Filter explicit results** setting enables safe search functionality. When enabled, the setting is stored in a browser cookie as `FF=` and passed to the server each time the user performs a Google search.<br><br>Appending `safe=active` to a Google search query URL also enables the strictest safe search settings.<br><br>**Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address**<br><br>Alternatively, Google provides servers that Lock SafeSearch (forcesafesearch.google.com) settings in every Google and YouTube search. By adding a DNS entry for `www.google.com` and `www.youtube.com` (and other relevant Google and YouTube country subdomains) that includes a CNAME record pointing to `forcesafesearch.google.com` to your DNS server configuration, you can ensure that all users on your network are using strict safe search settings every time they perform a Google or YouTube search. Keep in mind, however, that this solution is not compatible with Safe Search Enforcement on the firewall. Therefore, if you are using this option to force safe search on Google, the best practice is to block access to other search engines on the firewall by creating custom URL categories and adding them to the block list in the URL filtering profile.<br><br> If you plan to use the Google Lock SafeSearch solution, consider configuring DNS Proxy (**Network > DNS Proxy**) and setting the inheritance source as the Layer 3 interface on which the firewall receives DNS settings from service provider via DHCP. You would configure the DNS proxy with **Static Entries** for www.google.com and www.youtube.com, using the local IP address for the forcesafesearch.google.com server. |
| Yahoo | Offers safe search on individual computers only. The Yahoo Search Preferences include three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `vm=` and passed to the server each time the user performs a Yahoo search.<br><br>Appending `vm=r` to a Yahoo search query URL also enables the strictest safe search settings.<br><br> When performing a search on Yahoo Japan (yahoo.co.jp) while logged into a Yahoo account, end users must also enable the **SafeSearch Lock** option. |
| Bing | Offers safe search on individual computers or through their Bing in the Classroom program. The Bing Settings include three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `adlt=` and passed to the server each time the user performs a Bing search.<br><br>Appending `adlt=strict` to a Bing search query URL also enables the strictest safe search settings.<br><br>The Bing SSL search engine does not enforce the safe search URL parameters and you should therefore should consider blocking Bing over SSL for full safe search enforcement. |

 © Palo Alto Networks, Inc.

## Container Pages

A container page is the main page that a user accesses when visiting a website, but additional websites may be loaded within the main page. If the **Log Container page only** option is enabled in the URL filtering profile, only the main container page will be logged, not subsequent pages that may be loaded within the container page. Because URL filtering can potentially generate a lot of log entries, you may want to turn on this option, so log entries will only contain those URIs where the requested page file name matches the specific mime-types. The default set includes the following mime-types:

- application/pdf
- application/soap+xml
- application/xhtml+xml
- text/html
- text/plain
- text/xml

If you have enabled the **Log container page only** option, there may not always be a correlated URL log entry for threats detected by Antivirus or Vulnerability Protection.

## HTTP Header Logging

URL filtering provides visibility and control over web traffic on your network. For improved visibility into web content, you can configure the URL Filtering profile to log HTTP header attributes included in a web request. When a client requests a web page, the HTTP header includes the user agent, referer, and x-forwarded-for fields as attribute-value pairs and forwards them to the web server. When enabled for logging HTTP headers, the firewall logs the following attribute-value pairs in the URL Filtering logs:

| Attribute | Description |
|---|---|
| User-Agent | The web browser that the user used to access the URL, for example, Internet Explorer. This information is sent in the HTTP request to the server. |
| Referer | The URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested. |
| X-Forwarded-For | The header field option that preserves the IP address of the user who requested the web page. It allows you to identify the IP address of the user particularly if you have a proxy server on your network, where all requests might seem to originate from the proxy server's IP address. |

 © Palo Alto Networks, Inc.

## URL Filtering Response Pages

The firewall provides three predefined response pages that display by default when a user attempts to browse to a site in a category that is configured with one of the block actions in the URL Filtering Profile (block, continue, or override) or when Safe Search Enforcement is enabled:

- **URL Filtering and Category Match Block Page**—Access blocked by a URL filtering profile or because the URL category is blocked by a security policy.

- **URL Filtering Continue and Override Page**—Page with initial block policy that allows users to bypass the block. With the override page, after clicking **Continue**, the user must supply a password to override the policy that blocks the URL.

- **URL Filtering Safe Search Block Page**—Access blocked by a security policy with a URL filtering profile that has the Safe Search Enforcement option enabled. The user will see this page if a search is performed using Google, Bing, Yahoo, or Yandex and their browser or search engine account setting for Safe Search is not set to strict.

You can either use the predefined pages, or you can Customize the URL Filtering Response Pages to communicate your specific acceptable use policies and/or corporate branding. In addition, you can use the URL Filtering Response Page Variables for substitution at the time of the block event or add one of the supported Response Page References to external images, sounds, or style sheets.

**URL Filtering Response Page Variables**

| Variable | Usage |
|---|---|
| <user/> | The firewall replaces the variable with the username (if available via User-ID) or IP address of the user when displaying the response page. |
| <url/> | The firewall replaces the variable with the requested URL when displaying the response page. |
| <category/> | The firewall replaces the variable with the URL filtering category of the blocked request. |
| <pan_form/> | HTML code for displaying the continue button on the URL Filtering Continue and Override page. |

You can also add code that triggers the firewall to display different messages using depending on what URL category the user is attempting to access. For example, the following code snippet from a response page specifies to display Message 1 if the URL category is games, Message 2 if the category as travel, or Message 3 if the category is kids:

```
var cat = "<category/>";
switch(cat)
{
   case 'games':
      document.getElementById("warningText").innerHTML = "Message 1";
      break;
   case 'travel':
      document.getElementById("warningText").innerHTML = "Message 2";
      break;
   case 'kids':
      document.getElementById("warningText").innerHTML = "Message 3";
      break;
}
```

Only a single HTML page can be loaded into each virtual system for each type of block page. However, other resources such as images, sounds, and cascading style sheets (CSS files) can be loaded from other servers at time the response page is displayed in the browser. All references must include a fully qualified URL.

**Response Page References**

| Reference Type | Example HTML Code |
|---|---|
| Image | `<img src="http://virginiadot.org/images/Stop-Sign-gif.gif">` |
| Sound | `<embed src="http://simplythebest.net/sounds/WAV/WAV_files/ movie_WAV_files/ do_not_go.wav" volume="100" hidden="true" autostart="true">` |
| Style Sheet | `<link href="http://example.com/style.css" rel="stylesheet" type="text/css" />` |
| Hyperlink | `<a href="http://en.wikipedia.org/wiki/Acceptable_use_policy">View Corporate Policy</a>` |

 © Palo Alto Networks, Inc.

## URL Category as Policy Match Criteria

URL categories can be used as a match criteria in a policy to provide more granularity in the policy. For example, you may have a decryption policy defined, but you would like specific websites to bypass decryption. To do this, you would configure a decryption policy with the *no-decrypt* action and a URL category would be defined as match criteria for the policy rule, so the policy would only match traffic flows to websites that are part of the specified category.

The following table describes the policy types that can utilize URL categories:

| Policy Type | Description |
|---|---|
| Captive Portal | To ensure that users authenticate before being allowed access to a specific category, you can attach a URL category as a match criterion for the captive portal policy. |
| Decryption | Decryption policies can use URL categories as match criteria to determine if specified websites should be decrypted or not. For example, if you have a decryption policy with the action decrypt for all traffic between two zones, there may be specific website categories, such as *financial-services* and/or *health-and-medicine*, that should not be decrypted. In this case, you would create a new decryption policy with the action of *no-decrypt* that precedes the decrypt policy and then defines a list of URL categories as match criteria for the policy. By doing this, each URL category that is part of the no-decrypt policy will not be decrypted. You could also configure a custom URL category to define your own list of URLs that can then be used in the no-decrypt policy. |
| QoS | A QoS policy can use URL categories to determine throughput levels for specific website categories. For example, you may want to allow the streaming-media category, but limit throughput by adding the URL category as match criteria to the QoS policy. |
| Security | In security policies you can use URL categories both as a match criteria in the **Service/URL Category** tab, and in URL filtering profiles that are attached in the **Actions** tab.<br>If for example, the IT-security group in your company needs access to the hacking category, while all other users are denied access to the category, you must create the following rules:<br>• A security rule that allows the IT-Security group to access content categorized as hacking. The security rule references the *hacking* category in the S**ervices/URL Category** tab and IT-Security group in the **Users** tab.<br>• Another security rule that allows general web access for all users. To this rule you attach a URL filtering profile that blocks the hacking category.<br>The policy that allows access to hacking must be listed before the policy that blocks hacking. This is because security policy rules are evaluated top down, so when a user who is part of the security group attempts to access a hacking site, the policy rule that allows access is evaluated first and will allow the user access to the hacking sites. Users from all other groups are evaluated against the general web access rule which blocks access to the hacking sites. |

# PAN-DB Categorization Workflow

This section describes the PAN-DB components and describes the URL categorization resolution workflow that occurs as users access various URLs through the firewall.

▲ PAN-DB URL Categorization Components
▲ PAN-DB URL Categorization Workflow

## PAN-DB URL Categorization Components

The following table describes the PAN-DB components in detail. The BrightCloud system works similarly, but does not use an initial seed database.

| Component | Description |
|---|---|
| URL Filtering Seed Database | The initial seed database downloaded to the firewall is a small subset of the database that is maintained on the Palo Alto Networks URL cloud servers. The reason this is done is because the full database contains millions of URLs and many of these URLs may never be accessed by your users. When downloading the initial seed database, a region is selected (North America, Europe, APAC, Japan) and each region contains a subset of URLs most accessed for the given region. By doing this, the firewall will store a much smaller URL database, which greatly improves lookup performance. If a user accesses a website that is not in the local URL database, the full cloud database is queried and the firewall will then add the new URL to the local database. In other words, the local database on the firewall will be continually populated/customized based on user activity. Note that the local customized URL database will be cleared if PAN-DB seed database is re-downloaded or if you change the URL database vendor from PAN-DB to BrightCloud. |
| Cloud Service | The PAN-DB cloud service is implemented using Amazon Web Services (AWS). AWS provides a distributed high-performance and stable environment for seed database downloads and URL lookups for Palo Alto Networks firewalls and communication is performed over SSL. The AWS cloud systems hold the entire PAN-DB and is updated as new URLs are identified. The PAN-DB cloud service supports an automated mechanism to update the firewall's local URL database if the version does not match. Each time the firewall queries the cloud servers for URL lookups, it will also check for critical updates. If there have been no queries to the cloud servers for more than 30 minutes, the firewall will check for updates on the cloud systems.<br>The cloud system also provides a mechanism to submit URL category change requests. This is performed through the test-a-site service and is available directly from the device (URL filtering profile setup) and from the Palo Alto Networks Test A Site website. You can also submit a URL categorization change request directly from the URL filtering log on the firewall in the log details section. |

 © Palo Alto Networks, Inc.

| Component | Description |
|---|---|
| Management Plane (MP) URL Cache | When PAN-DB is activated on the firewall, a seed database is downloaded from one of the PAN-DB cloud servers to initially populate the local cache, which is done to improve lookup performance. Each regional seed database contains the top URLs for the region and the size of the seed database (number of URL entries) also depends on the device platform. The URL MP cache is automatically written to the firewall's local drive every eight hours, before the firewall is rebooted, or when the cloud upgrades the URL database version on the firewall. After rebooting the firewall, the file that was saved to the local drive will be loaded to the MP cache. A least recently used (LRU) mechanism is also implemented in the URL MP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs. |
| Dataplane (DP) URL Cache | This is a subset of the MP cache and is a customized, dynamic URL database that is stored in the dataplane (DP) and is used to improve URL lookup performance. The URL DP cache is cleared at each firewall reboot. The number of URLs that are stored in the URL DP cache varies by hardware platform and the current URLs stored in the TRIE (data structure). A least recently used (LRU) mechanism is implemented in the DP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs. Entries in the URL DP cache expire after a specified period of time and the expiration period cannot be changed by the administrator. |

## PAN-DB URL Categorization Workflow

When a user attempts to access a URL and the URL category needs to be determined, the firewall will compare the URL with the following components (in order) until a match has been found:

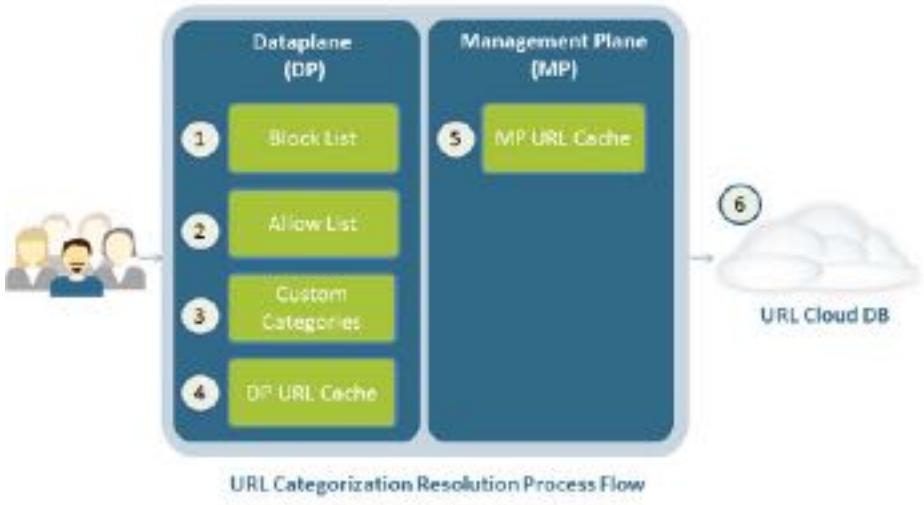


URL Categorization Resolution Process Flow

If a URL query matches an expired entry in the URL DP cache, the cache responds with the expired category, but also sends a URL categorization query to the management plane. This is done to avoid unnecessary delays in the DP, assuming that the frequency of changing categories is low. Similarly, in the URL MP cache, if a URL query from the DP matches an expired entry in the MP, the MP responds to the DP with the expired category and will also send a URL categorization request to the cloud service. Upon getting the response from the cloud, the firewall will resend the updated response to the DP.

As new URLs and categories are defined or if critical updates are needed, the cloud database will be updated. Each time the firewall queries the cloud for a URL lookup or if no cloud lookups have occurred for 30 minute, the database versions on the firewall be compared and if they do not match, an incremental update will be performed.

 © Palo Alto Networks, Inc.

# Enable a URL Filtering Vendor

To enable URL filtering on a firewall, you must purchase and activate a URL Filtering license for one of the supported URL Filtering Vendors and then install the database. Use one of the following procedures to enable a URL filtering license on the firewall set up the database. Only one URL filtering license can be active on a firewall.

Starting with PAN-OS 6.0, firewalls managed by Panorama do not need to be running the same URL filtering vendor that is configured on Panorama. For firewalls running PAN-OS 6.0 or later, when a mismatch is detected between the vendor enabled on the firewalls and what is enabled on Panorama, the firewalls can automatically migrate URL categories and/or URL profiles to (one or more) categories that align with that of the vendor enabled on it. For guidance on how to configure URL filtering Panorama if you are managing firewalls running different PAN-OS versions, refer to the Panorama Administrator's Guide.

If you have valid licenses for both PAN-DB and BrightCloud, activating the PAN-DB license automatically deactivates the BrightCloud license (and vice versa).

- Enable PAN-DB URL Filtering
- Enable BrightCloud URL Filtering

# Enable PAN-DB URL Filtering

| Enable PAN-DB URL Filtering | |
|---|---|
| Step 1 Obtain and install a PAN-DB URL filtering license and confirm that it is installed.<br><br>If the license expires, PAN-DB URL filtering continues to work based on the URL category information that exists in the dataplane and management plane caches. However, URL cloud lookups and other cloud-based updates will not function until you install a valid license. | 1. Select **Device > Licenses** and, in the **License Management** section, select the license installation method:<br>• **Activate feature using authorization code**<br>• **Retrieve license keys from license server**<br>• **Manually upload license key**<br>2. After installing the license, confirm that the PAN-DB URL Filtering section, **Date Expires** field, displays a valid date. |
| Step 2 Download the initial seed database and activate PAN-DB URL filtering.<br><br>The firewall must have Internet access; you cannot manually upload the PAN-DB. | 1. In the PAN-DB URL Filtering section, **Download Status** field, click **Download Now**.<br>2. Choose a region (North America, Europe, APAC, Japan) and then click **OK** to start the download.<br>3. After the download completes, click **Activate**.<br><br>If PAN-DB is already the active URL filtering vendor and you click **Re-Download**, this will reactivate PAN-DB by clearing the dataplane and management plane caches and replacing them with the contents of the new seed database. You should avoid doing this unless it is necessary, as you will lose your cache, which is customized based on the web traffic that has previously passed through the firewall based on user activity. |
| Step 3 Schedule the firewall to download dynamic updates for Applications and Threats.<br><br>A Threat Prevention license is required to receive content updates, which covers Antivirus and Applications and Threats. | 1. Select **Device > Dynamic Updates**.<br>2. In the Schedule field in the Applications and Threats section, click the **None** link to schedule periodic updates.<br><br>You can only schedule dynamic updates if the firewall has direct Internet access.If updates are already scheduled in a section, the link text displays the schedule settings.<br><br>The Applications and Threats updates might contain updates for URL filtering related to the **Safe Search Enforcement** option in the URL filtering profile (**Objects > Security Profiles > URL Filtering**). For example, if Palo Alto Networks adds support for a new search provider vendor or if the method used to detect the Safe Search setting for an existing vendor changes, the Application and Threats updates will include that update. |


© Palo Alto Networks, Inc.

# Enable BrightCloud URL Filtering

| Enable BrightCloud URL Filtering | | |
|---|---|---|
| Step 1 | Obtain and install a BrightCloud URL filtering license and confirm that it is installed.<br><br>BrightCloud has an option in the URL filtering profile (**Objects > Security Profiles > URL Filtering**) to either allow all categories or block all categories if the license expires. | 1. Select **Device > Licenses** and, in the **License Management** section, select the license installation method:<br>• **Activate feature using authorization code**<br>• **Retrieve license keys from license server**<br>• **Manually upload license key**<br>2. After installing the license, confirm that the BrightCloud URL Filtering section, **Date Expires** field, displays a valid date. |
| Step 2 | Install the BrightCloud database.<br><br>The way you do this depends on whether or not the firewall has direct Internet access. | **Firewall with Direct Internet Access**<br>In the **Device > Licenses** page, BrightCloud URL Filtering section, Active field, click the **Activate** link to install the BrightCloud database. This operation automatically initiates a system reset.<br>**Firewall without Direct Internet Access**<br>1. Download the BrightCloud database to a host that has Internet access. The firewall must have access to the host:<br>a. On a host with Internet access, go to the Palo Alto Support website (https://support.paloaltonetworks.com) and log in.<br>b. In the Resources section, click **Dynamic Updates**.<br>c. In the BrightCloud Database section, click **Download** and save the file to the host.<br>2. Upload the database to the firewall:<br>a. Log in to the firewall, select **Device > Dynamic Updates** and click **Upload**.<br>b. For the **Type**, select **URL Filtering**.<br>c. Enter the path to the **File** on the host or click **Browse** to find it, then click **OK**. When the Status is *Completed*, click **Close**.<br>3. Install the database:<br>a. In the **Device > Dynamic Updates** page, click **Install From File**.<br>b. For the **Type**, select **URL Filtering**. The firewall automatically selects the file you just uploaded.<br>c. Click **OK** and, when the Result is *Succeeded*, click **Close**. |
| Step 3 | Enable cloud lookups for dynamically categorizing a URL if the category is not available on the local BrightCloud database. | 1. Access the firewall CLI.<br>2. Enter the following commands to enable dynamic URL filtering:<br>`configure`<br>`set deviceconfig setting url dynamic-url yes`<br>`commit` |

| Enable BrightCloud URL Filtering (Continued) | | |
|---|---|---|
| Step 4 | Schedule the firewall to download dynamic updates for Applications and Threats signatures and URL filtering.<br><br>You can only schedule dynamic updates if the firewall has direct Internet access.<br><br>The Applications and Threats updates might contain updates for URL filtering related to the **Safe Search Enforcement** option in the URL filtering profile. For example, if Palo Alto Networks adds support for a new search provider vendor or if the method used to detect the Safe Search setting for an existing vendor changes, the Application and Threats updates will include that update.<br><br>BrightCloud updates include a database of approximately 20 million websites that are stored on the firewall drive. You must schedule the URL filtering updates to receive database updates.<br><br> A Threat Prevention license is required to receive content updates, which covers Antivirus and Applications and Threats. | 1. Select **Device > Dynamic Updates**.<br>2. In the Applications and Threats section, Schedule field, click the **None** link to schedule periodic updates.<br>3. In the URL Filtering section, Schedule field, click the **None** link to schedule periodic updates.<br><br> If updates are already scheduled in a section, the link text displays the schedule settings. |

 © Palo Alto Networks, Inc.

# Determine URL Filtering Policy Requirements

The recommended practice for deploying URL filtering in your organization is to first start with a passive URL filtering profile that will alert on most categories. After setting the alert action, you can then monitor user web activity for a few days to determine patterns in web traffic. After doing so, you can then make decisions on the websites and website categories that should be controlled.

In the procedure that follows, threat-prone sites will be set to block and the other categories will be set to alert, which will cause all websites traffic to be logged. This may potentially create a large amount of log files, so it is best to do this for initial monitoring purposes to determine the types of websites your users are accessing. After determining the categories that your company approves of, those categories should then be set to allow, which will not generate logs. You can also reduce URL filtering logs by enabling the **Log container page only** option in the URL Filtering profile, so only the main page that matches the category will be logged, not subsequent pages/categories that may be loaded within the container page.

| **Configure and Apply a Passive URL Filtering Profile** | | |
|---|---|---|
| Step 1 | Create a new URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select the default profile and then click **Clone**. The new profile will be named **default-1**.<br>3. Select the **default-1** profile and rename it. For example, rename it to URL-Monitoring. |
| Step 2 | Configure the action for all categories to **alert**, except for threat-prone categories, which should remain blocked.<br> To select all items in the category list from a Windows system, click the first category, then hold down the shift key and click the last category—this will select all categories. Hold the control key (ctrl) down and click items that should be deselected. On a Mac, do the same using the shift and command keys. You could also just set all categories to alert and manually change the recommended categories back to block. | 1. In the section that lists all URL categories, select all categories.<br>2. To the right of the *Action* column heading, mouse over and select the down arrow and then select **Set Selected Actions** and choose **alert**.<br> 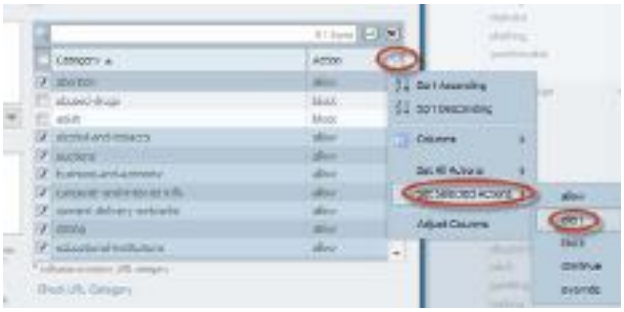 <br>3. To ensure that you block access to threat-prone sites, select the following categories and then set the action to **block**: abused-drugs, adult, gambling, hacking, malware. phishing, questionable, weapons.<br>4. Click **OK** to save the profile. |
| Step 3 | Apply the URL Filtering profile to the security policy rule(s) that allows web traffic for users. | 1. Select **Policies > Security** and select the appropriate security policy to modify it.<br>2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the new profile.<br>3. Click **OK** to save. |
| Step 4 | Save the configuration. | Click **Commit**. |

| Configure and Apply a Passive URL Filtering Profile (Continued) | | |
|---|---|---|
| Step 5 | View the URL filtering logs to determine all of the website categories that your users are accessing. In this example, some categories are set to block, so those categories will also appear in the logs.<br><br>For information on viewing the logs and generating reports, see Monitor Web Activity. | Select **Monitor > Logs > URL Filtering**. A log entry will be created for any website that exists in the URL filtering database that is in a category that is set to any action other than **allow**. |


© Palo Alto Networks, Inc.

# Monitor Web Activity

URL filtering logs and reports show all user web activity for URL categories that are set to **alert**, **block**, **continue**, or **override**. By monitoring the logs, you can gain a better understanding of the web activity of your user base to determine a web access policy.

The following topics describe how to monitor web activity:

- Interpret the URL Filtering Logs
- Use the ACC to Monitor Web Activity
- View URL Filtering Reports
- Configure Custom URL Filtering Reports

## Interpret the URL Filtering Logs

The following bullet points show examples of the URL filtering logs (**Monitor > Logs > URL filtering**).

- **Alert log**—In this log, the category is shopping and the action is alert.

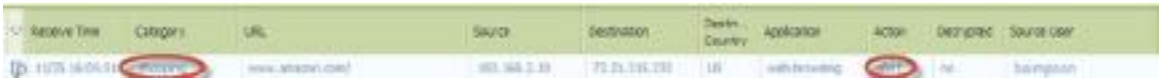

- **Block log**—In this log, the category alcohol-and-tobacco was set to block, so the action is block-url and the user will see a response page indicating that the website was blocked.

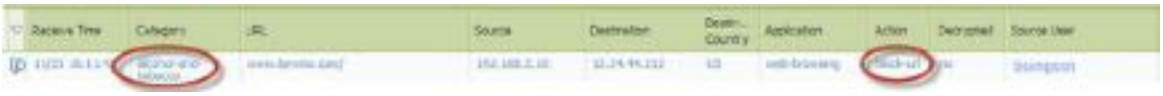

- **Alert log on encrypted website**—In this example, the category is social-networking and the application is facebook-base, which is required to access the Facebook website and other Facebook applications. Because faceboook.com is always encrypted using SSL, the traffic was decrypted by the firewall, which allows the website to be recognized and controlled if needed.

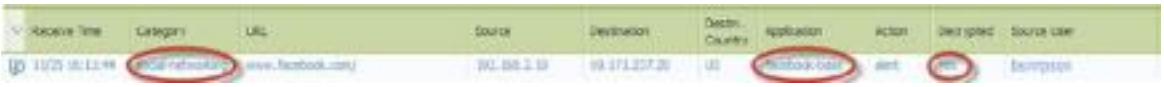

You can also add several other columns to your URL Filtering log view, such as: to and from zone, content type, and whether or not a packet capture was performed. To modify what columns to display, click the down arrow in any column and select the attribute to display.

 © Palo Alto Networks, Inc.

To view the complete log details and/or request a category change for the given URL that was accessed, click the log details icon in the first column of the log.

## Use the ACC to Monitor Web Activity

For a quick view of the most common categories being accessed in your environment, select the **ACC** tab and scroll down to the **URL Filtering** section. Along the top of this window, you can also set the time range, sort by option, and define how many results will appear. Here you will see the most popular categories that are accessed by your users, sorted by the most popular at the top of the list. In this example, computer-and-internet-info is the most accessed category, followed by private-ip-addresses (internal servers), and search-engines. In the drop-down in the upper right of the statistics, you can also choose to list by URL Categories, Blocked URL Categories, and Blocked URLs.

 © Palo Alto Networks, Inc.

## View URL Filtering Reports

To view the default URL filtering reports, select **Monitor > Reports** and under the **URL Filtering Reports** section, choose one of the reports. You can generate reports on URL Categories, URL users, Web Sites accessed, Blocked Categories, and more. The reports are based on a 24-hour period and the day is selected by choosing a day in the calendar section. You can also export the report to PDF, CSV, or XML.

## View the User Activity Report

This report provides a quick method of viewing user or group activity and also provides an option to view browse time activity.

| Generate a User Activity Report | |
|---|---|
| Step 1 Configure a User Activity Report | 1. Select **Monitor > PDF Reports > User Activity Report**.<br>2. Enter a report **Name** and select the report type. Select **User** to generate a report for one person, or select **Group** for a group of users.<br>You must Enable User-ID in order to be able to select user or group names. If User-ID is not configured, you can select the type **User** and enter the IP address of the user's computer.<br>3. Enter the Username/IP address for a user report or enter the group name for a user group report.<br>4. Select the time period. You can select an existing time period, or select **Custom**.<br>5. Select the **Include Detailed Browsing** check box, so browsing information is included in the report.<br> |
| Step 2 Run the user activity report and then download the report. | 1. Click **Run Now**.<br>2. After the report is generated, click the **Download User Activity Report** link.<br> 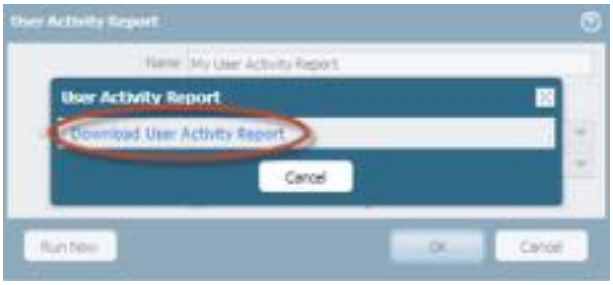 <br>3. After the report is downloaded, click **Cancel** and then click **OK** to save the report. |

 © Palo Alto Networks, Inc.

**Generate a User Activity Report (Continued)**

Step 3 View the user activity report by opening the PDF file that was downloaded. The top of the report will contain a table of contents similar to the following:

User Activity Report for 192.168.2.10
Tuesday, December 31, 2013 09:35:47 - 10:35:46


Application Usage 2
Traffic Summary by URL Category 3
Browsing Summary by URL Category 4
Browsing Summary by Website 5
Blocked Browsing Summary by Website 6
Detailed Web Browsing Activity 7


Step 4 Click an item in the table of contents to view details. For example, click *Traffic Summary by URL Category* to view statistics for the selected user or group.

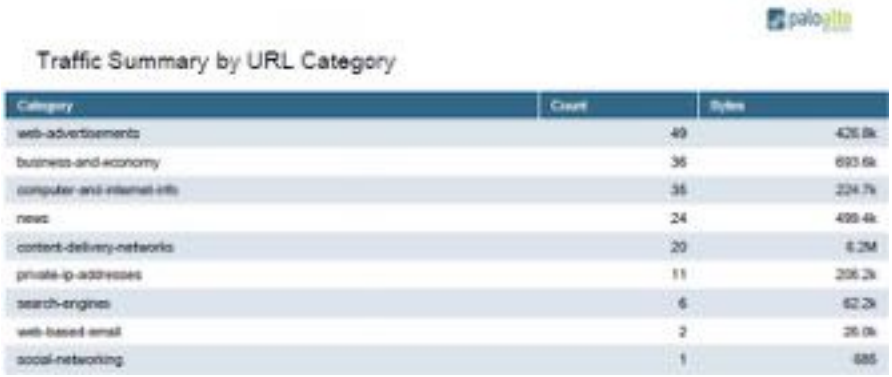
Traffic Summary by URL Category

| Category | Count | Bytes |
|---|---|---|
| web-advertisements | 49 | 426.8k |
| business-and-economy | 36 | 693.6k |
| computer-and-internet-info | 35 | 224.7k |
| news | 24 | 499.4k |
| content-delivery-networks | 20 | 6.2M |
| private-ip-addresses | 11 | 206.2k |
| search-engines | 6 | 62.2k |
| web-based-email | 2 | 26.0k |
| social-networking | 1 | 686 |

## Configure Custom URL Filtering Reports

To generate a detailed report that can also be scheduled, you can configure a custom report and select from a list of all available URL filtering log fields.

| Configure a Custom URL Filtering Report | | |
|---|---|---|
| Step 1 | Add a new custom report. | 1. Select **Monitor > Manage Custom Reports** and click **Add**.<br>2. Enter a report **Name**, for example, My-URL-Custom-Report.<br>3. From the **Database** drop-down, select **URL Log**. |
| Step 2 | Configure report options. | 1. Select the **Time Frame** drop-down and select a range.<br>2. (Optional) To customize how the report is sorted and grouped, select **Sort By** and chose the number of items to display (top 25 for example) and then select **Group By** and select an option such as **Category**, and then select how many groups will be defined.<br>3. In the **Available Columns** list, select the fields to include the report. The following columns are typically used for a URL report:<br>• Action<br>• Category<br>• Destination Country<br>• Source User<br>• URL<br> |
| Step 3 | Run the report to check the results. If the results are satisfactory, set a schedule to run the report automatically. | 1. Click the **Run Now** icon to immediately generate the report that will appear in a new tab.<br>2. (Optional) Click the **Schedule** check box to run the report once per day. This will generate a daily report that details web activity over the last 24 hours. To access the report, select **Monitor > Report** and then expand **Custom Reports** on the right column and select the report. |
| Step 4 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

# Configure URL Filtering

After you Determine URL Filtering Policy Requirements, you should have a basic understanding of what types of websites and website categories your users are accessing. With this information, you are now ready to create custom URL filtering profiles and attach them to the security policy rule(s) that allow web access.

| Configure Website Controls | | |
|---|---|---|
| Step 1 | Create a URL Filtering profile or select an existing one.<br> <br>Because the default URL filtering profile blocks risky and threat-prone content, it is a best practice to clone this profile rather than creating a new profile to preserve these default settings. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>Select the default profile and then click **Clone**. The new profile will be named default-1.<br>2. Select the new profile and rename it. |
| Step 2 | Define how to control access to web content. | In the **Categories** tab, for each category that you want visibility into or control over, select a value from the **Action** column as follows:<br>• If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **allow**.<br>• For visibility into traffic to sites in a category, select **alert**.<br>• To deny access to traffic that matches the category and to enable logging of the blocked traffic, select **block**.<br>• To require users to click Continue to proceed to a questionable site, select **continue**.<br>• To only allow access if users provide a configured password, select override. For more details on this setting, see Configure URL Admin Override. |
| Step 3 | Define websites that should always be blocked or allowed.<br>For example, to reduce URL filtering logs, you may want add you corporate websites in the allow list, so no logs will be generated for those sites. Or, if there is a website this is being overly used and is not work related in any way, you can add it to the block list.<br>Items in the block list will always be blocked regardless of the action for the associated category, and URLs in the allow list will always be allowed.<br>For more information on the proper format and wildcards usage, see Block and Allow Lists. | 1. In the URL filtering profile, enter URLs or IP addresses in the **Block List** and select an action:<br>• **block**—Block the URL.<br>• **continue**—Prompt users click **Continue** to proceed to the web page.<br>• **override**—The user will be a prompted for a password to continue to the website.<br>• **alert**—Allow the user to access the website and add an alert log entry in the URL log.<br>2. For the **Allow** list, enter IP addresses or URLs that should always be allowed. Each row must be separated by a new line.<br>3. (Optional) Enable Safe Search Enforcement. |

| Configure Website Controls | | |
|---|---|---|
| Step 4 | Modify the setting to log Container Pages only. | The **Log container page only** option is enabled by default so that only the main page that matches the category is logged, not subsequent pages/categories that may be loaded within the container page. To enable logging for all pages/categories, clear the **Log container page only** check box. |
| Step 5 | Enable HTTP Header Logging for one or more of the supported HTTP header fields. | To log an HTTP header field, select one or more of the following fields to log:<br>• **User-Agent**<br>• **Referer**<br>• **X-Forwarded-For** |
| Step 6 | Save the URL filtering profile. | Click **OK**. |
| Step 7 | (Optional) Enable response pages on the ingress interface (the interface that first receives traffic for your users).<br>This option is required if you enable the **continue** or **override** action for any URL category.<br> You can also Customize the URL Filtering Response Pages. | 1. Select **Network > Network Profiles > Interface Mgmt** and either add a new profile or edit an existing profile.<br>2. Click the **Response Pages** check box to enable.<br>3. Click **OK** to save the profile.<br>4. Select **Network > Interfaces** and then edit the Layer 3 interface or VLAN interface that is your ingress interface.<br>5. On the **Advanced** tab, select the Interface Mgmt profile that has the response page option enabled.<br>6. Click **OK** to save the interface configuration. |
| Step 8 | Save the configuration.<br> To test the URL filtering configuration, simply access a website in a category that is set to block or continue to see if the appropriate action is performed. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# Log Client IP Addresses for Proxied HTTP/HTTPS Requests

If you have a proxy server deployed between the users on your network and the firewall, the firewall might identify the proxy server IP address as the source IP address in the traffic that the proxy forwards rather than the IP address of the client that requested the content. In many cases, the proxy server adds an X-Forwarded-For (XFF) header to the traffic packets that includes the IP address of the source client. In such cases, if you want to identify the client that requested the content, you must configure the firewall to log the XFF values. This enables the firewall to get the client IP address, the username if available, the IP address of the last proxy server traversed in a proxy chain, or any string that the field stores. You can then use this information to monitor web access on your network.

You can configure the firewall to log XFF values in the following ways:

- Replace Source IP Addresses with XFF Values in URL Filtering Logs
- Add XFF Values to URL Filtering Logs

## Replace Source IP Addresses with XFF Values in URL Filtering Logs

You can configure the firewall to replace the value in the source IP address field of URL Filtering logs with the XFF values of web request packets. The XFF values can be the IP address of the client that originated the request or, in a proxy chain, the IP address of the last proxy server that processed the request. To ensure that attackers can't read and exploit the XFF values when packets exit the firewall to retrieve content from an external server, you can also configure the firewall to strip the XFF values from outgoing packets. These options are not mutually exclusive: if you configure both, the firewall zeroes out the XFF values only after logging them.

| **Replace Source IP Addresses with XFF Values in URL Filtering Logs** | | |
|---|---|---|
| Step 1 | Enable logging of XFF values. | 1. Select **Device > Setup > Content-ID** and edit the URL Filtering section.<br>2. Select the **x-forwarded-for** check box. |
| Step 2 | Remove XFF values from outgoing web requests. | 1. Select the **Strip x-forwarded-for** check box.<br>2. Click **OK** and **Commit**. |
| Step 3 | Verify the firewall is logging XFF values. | 1. Select **Monitor > Logs > URL Filtering**.<br>2. Verify that the IP addresses in the Source column reflect the XFF values. |

## Add XFF Values to URL Filtering Logs

You can configure the firewall to add the XFF field values of web requests to URL Filtering logs. The XFF value can be a client IP address, a username if available, or any value of up to 128 characters that the field stores.

This method of logging XFF values doesn't replace the source IP address values in URL Filtering logs. When you view the logs, the source field displays the IP address of the proxy server if one is deployed between the user clients and the firewall.

| Add XFF Values to URL Filtering Logs | | |
|---|---|---|
| Step 1 | Configure a URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select an existing profile or **Add** a new profile and enter a descriptive **Name**.<br>You can't enable XFF logging in the default URL Filtering profile.<br>3. In the **Categories** tab, define how to control access to web content. For each category that you want visibility into or control over, select a value in the **Action** column:<br>• If you don't care about traffic to a particular category (that is you neither want to block it nor log it), select **Allow**.<br>• For visibility into traffic to sites in a category, select **Alert**.<br>• To deny access to traffic that matches the category and to enable logging of the blocked traffic, select **Block**.<br>4. Select the **Settings** tab and select the **X-Forwarded-For** check box.<br>5. Click **OK** to save the profile. |
| Step 2 | Attach the URL Filtering profile to a policy rule. | 1. Select **Policies > Security** and click the rule.<br>2. Select the **Actions** tab, set the **Profile Type** to **Profiles**, and select the **URL Filtering** profile you just created.<br>3. Click **OK** and **Commit**. |
| Step 3 | Verify the firewall is logging XFF values. | 1. Select **Monitor > Logs > URL Filtering**.<br>2. Display the XFF values in one of the following ways:<br>• To display the XFF value for a single log—Click the icon for the log to displays its details. The HTTP Headers section displays the X-Forwarded-For value.<br>• To display the XFF values for all logs—Open the drop-down in any column header, select **Columns**, and select the **X-Forwarded-For** check box. The page then displays an X-Forwarded-For column. |


© Palo Alto Networks, Inc.

# Customize the URL Filtering Response Pages

The firewall provides three predefined URL Filtering Response Pages that display by default when a user attempts to browse to a site in a category that is configured with one of the block actions in the URL Filtering Profile (block, continue, or override) or when Safe Search Enforcement blocks a search attempt. However, you can create your own custom response pages with your corporate branding, acceptable use policies, links to your internal resources as follows:

| **Customize the URL Filtering Response Pages** | | |
|---|---|---|
| Step 1 | Export the default response page(s). | 1. Select **Device > Response Pages**.<br>2. Select the link for the URL filtering response page you want to modify.<br>3. Click the response page (predefined or shared) and then click the **Export** link and save the file to your desktop. |
| Step 2 | Edit the exported page. | 1. Using the HTML text editor of your choice, edit the page:<br>• If you want the response page to display custom information about the specific user, URL, or category that was blocked, add one or more of the supported URL Filtering Response Page Variables.<br>• If you want to include custom images (such as your corporate logo), a sound, or style sheet, or link to another URL, for example to a document detailing your acceptable web use policy, include one or more of the supported Response Page References.<br>2. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding. For example, in Notepad you would select **UTF-8** from the **Encoding** drop-down in the Save As dialog. |
| Step 3 | Import the customized response page. | 1. Select **Device > Response Pages**.<br>2. Select the link that corresponds to the URL Filtering response page you edited.<br>3. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.<br>4. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select **shared** to make it available to all virtual systems.<br>5. Click **OK** to import the file. |
| Step 4 | Save the new response page(s). | **Commit** the changes. |
| Step 5 | Verify that the new response page displays. | From a browser, go to the URL that will trigger the response page. For example, to see a modified URL Filtering and Category Match response page, browse to URL that your URL filtering policy is set to block. |

# Configure URL Admin Override

In some cases there may be URL categories that you want to block, but allow certain individuals to browse to on occasion. In this case, you would set the category action to **override** and define a URL admin override password in the firewall Content-ID configuration. When users attempt to browse to the category, they will be required to provide the override password before they are allowed access to the site. Use the following procedure to configure URL admin override:

| **Configure URL Admin Override** | |
|---|---|
| Step 1 Set the URL admin override password. | 1. Select **Device > Setup > Content ID**.<br>2. In the **URL Admin Override** section, click **Add**.<br>3. In the **Location** field, select the virtual system to which this password applies.<br>4. Enter the **Password** and **Confirm Password**.<br>5. Select a **Server Certificate** for the firewall to present to the user if the site with the override is an HTTPS site.<br>6. Select the **Mode** for prompting the user for the password:<br>• **Transparent**—The firewall intercepts the browser traffic destined for site in a URL category you have set to override and impersonates the original destination URL, issuing an HTTP 401 to prompt for the password. Note that the client browser will display certificate errors if it does not trust the certificate.<br>• **Redirect**—The firewall intercepts HTTP or HTTPS traffic to a URL category set to override and redirects the request to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to prompt for the override password. If you select this option, you must provide the **Address** (IP address or DNS hostname) to which to redirect the traffic.<br>7. Click **OK**. |
| Step 2 (Optional) Set a custom override period. | 1. Edit the URL Filtering section.<br>2. To change the amount of time users can browse to a site in a category for which they have successfully entered the override password, enter a new value in the **URL Admin Override Timeout** field. By default, users can access sites within the category for 15 minutes without re-entering the password.<br>3. To change the amount of time users are blocked from accessing a site set to override after three failed attempts to enter the override password, enter a new value in the **URL Admin Lockout Timeout** field. By default, users are blocked for 30 minutes.<br>4. Click **OK**. |


© Palo Alto Networks, Inc.

| Configure URL Admin Override (Continued) | | |
|---|---|---|
| Step 3 | (Redirect mode only) Create a Layer 3 interface to which to redirect web requests to sites in a category configured for override. | 1. Create a management profile to enable the interface to display the URL Filtering Continue and Override Page response page:<br>a. Select **Network > Interface Mgmt** and click **Add**.<br>b. Enter a **Name** for the profile, select **Response Pages**, and then click **OK**.<br>2. Create the Layer 3 interface. Be sure to attach the management profile you just created (on the **Advanced > Other Info** tab of the Ethernet Interface dialog). |
| Step 4 | (Redirect mode only) To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface to which you are redirecting web requests to a site in a URL category configured for override.You can either generate a self-signed certificate or import a certificate that is signed by an external CA. | To use a self-signed certificate, you must first create a root CA certificate and then use that CA to sign the certificate you will use for URL admin override as follows:<br>1. To create a root CA certificate, select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**. Enter a **Certificate Name**, such as RootCA. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed). Make sure you select the **Certificate Authority** check box and then click **Generate** the certificate.<br>2. To create the certificate to use for URL admin override, click **Generate**. Enter a **Certificate Name** and enter the DNS hostname or IP address of the interface as the **Common Name**. In the **Signed By** field, select the CA you created in the previous step. Add an IP address attribute and specify the IP address of the Layer 3 interface to which you will be redirecting web requests to URL categories that have the override action.<br>3. **Generate** the certificate.<br>4. To configure clients to trust the certificate, select the CA certificate on the **Device Certificates** tab and click **Export**. You must then import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO). |
| Step 5 | Specify which URL categories require an override password to enable access. | 1. Select **Objects > URL Filtering** and either select an existing URL filtering profile or **Add** a new one.<br>2. On the **Categories** tab, set the Action to **override** for each category that requires a password.<br>3. Complete any remaining sections on the URL filtering profile and then click **OK** to save the profile. |
| Step 6 | Apply the URL Filtering profile to the security policy rule(s) that allows access to the sites requiring password override for access. | 1. Select **Policies > Security** and select the appropriate security policy to modify it.<br>2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the profile.<br>3. Click **OK** to save. |
| Step 7 | Save the configuration. | Click **Commit**. |

# Enable Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos for search query return traffic. You can configure Safe Search Enforcement the Palo Alto Networks next-generation firewall to prevent search requests that do not have the strictest safe search settings enabled.

The Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address is not compatible with Safe Search Enforcement on the firewall.

There are two ways to enforce Safe Search on the firewall:

- Block Search Results that are not Using Strict Safe Search Settings
- Enable Transparent Safe Search Enforcement

## Block Search Results that are not Using Strict Safe Search Settings

By default, when you enable safe search enforcement, when a user attempts to perform a search without using the strictest safe search settings, the firewall will block the search query results and display the URL Filtering Safe Search Block Page. This page provides a link to the search settings page for the corresponding search provider so that the end user can enable the safe search settings. If you plan to use this default method for enforcing safe search, you should communicate the policy to your end users prior to deploying the policy. See Table: Search Provider Safe Search Settings for details on how each search provider implements safe search. The default URL Filtering Safe Search Block Page provides a link to the search settings for the corresponding search provider. You can optionally Customize the URL Filtering Response Pages.

Alternatively, to enable safe search enforcement so that it is transparent to your end users, configure the firewall to Enable Transparent Safe Search Enforcement.

 © Palo Alto Networks, Inc.

| Enable Safe Search Enforcement | | |
|---|---|---|
| Step 1 | Enable Safe Search Enforcement in the URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select an existing profile to modify, or clone the default profile to create a new profile.<br>3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.<br>4. (Optional) Restrict users to specific search engines:<br>a. On the **Categories** tab, set the **search-engines** category to **block**.<br>b. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:<br>`www.google.com`<br>`www.bing.com`<br>5. Configure other settings as necessary to:<br>• Define how to control access to web content.<br>• Define websites that should always be blocked or allowed.<br>6. Click **OK** to save the profile. |
| Step 2 | Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet. | 1. Select **Policies > Security** and select a rule to which to apply the URL filtering profile that you just enabled for Safe Search Enforcement.<br>2. On the **Actions** tab, select the **URL Filtering** profile.<br>3. Click **OK** to save the security policy rule. |
| Step 3 | Enable SSL Forward Proxy decryption.<br>Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings. | 1. Add a custom URL category for the search sites:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as SearchEngineDecryption.<br>c. **Add** the following to the Sites list:<br>`www.bing.*`<br>`www.google.*`<br>`search.yahoo.*`<br>d. Click **OK** to save the custom URL category object.<br>2. Follow the steps to Configure SSL Forward Proxy.<br>3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**. |

| Enable Safe Search Enforcement (Continued) | | |
|---|---|---|
| Step 4 | (Optional, but recommended) Block Bing search traffic running over SSL.<br><br>Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL. | 1. Add a custom URL category for Bing:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as EnableBingSafeSearch.<br>c. **Add** the following to the Sites list:<br>`www.bing.com/images/*`<br>`www.bing.com/videos/*`<br>d. Click **OK** to save the custom URL category object.<br>2. Create another URL filtering profile to block the custom category you just created:<br>a. Select **Objects > Security Profiles > URL Filtering**.<br>b. **Add** a new profile and give it a descriptive **Name**.<br>c. Locate the custom category in the Category list and set it to **block**.<br>d. Click **OK** to save the URL filtering profile.<br>3. Add a security policy rule to block Bing SSL traffic:<br>a. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.<br>b. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.<br>c. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.<br>d. Select **TCP** as the **Protocol** and set the **Destination Port** to `443`.<br>e. Click **OK** to save the rule.<br>f. Use the **Move** options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled. |
| Step 5 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

**Enable Safe Search Enforcement (Continued)**

| | |
|---|---|
| Step 6 Verify the Safe Search Enforcement configuration.<br><br>This verification step only works if you are using block pages to enforce safe search. If you are using transparent safe search enforcement, the firewall block page will invoke a URL rewrite with the safe search parameters in the query string. | 1. From a computer that is behind the firewall, disable the strict search settings for one of the supported search providers. For example, on bing.com, click the **Preferences** icon on the Bing menu bar.<br><br> 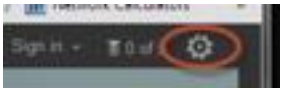 <br><br>2. Set the **SafeSearch** option to **Moderate** or **Off** and click **Save**.<br>3. Perform a Bing search and verify that the URL Filtering Safe Search Block page displays instead of the search results:<br><br> <br><br><br>4. Use the link in the block page to go to the search settings for the search provider and set the safe search setting back to the strictest setting (**Strict** in the case of Bing) and then click **Save**.<br>5. Perform a search again from Bing and verify that the filtered search results display instead of the block page. |

## Enable Transparent Safe Search Enforcement

If you want to enforce filtering of search query results with the strictest safe search filters, but you don't want your end users to have to manually configure the settings, you can enable transparent Safe Search Enforcement as follows. This functionality is supported on Google, Yahoo, and Bing search engines only and requires Content Release version 475 or later.

| **Enable Transparent Safe Search Enforcement** | | |
|---|---|---|
| Step 1 | Make sure the firewall is running Content Release version 475 or later. | 1. Select **Device > Dynamic Updates**.<br>2. Check the **Applications and Threats** section to determine what update is currently running.<br>3. If the firewall is not running the required update or later, click **Check Now** to retrieve a list of available updates.<br>4. Locate the required update and click **Download**.<br>5. After the download completes, click **Install**. |
| Step 1 | Enable Safe Search Enforcement in the URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select an existing profile to modify, or clone the default profile to create a new one.<br>3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.<br>4. (Optional) Allow access to specific search engines only:<br>a. On the **Categories** tab, set the **search-engines** category to **block**.<br>b. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:<br>`www.google.com`<br>`www.bing.com`<br>5. Configure other settings as necessary to:<br>• Define how to control access to web content.<br>• Define websites that should always be blocked or allowed.<br>6. Click **OK** to save the profile. |
| Step 2 | Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet. | 1. Select **Policies > Security** and select a rule to which to apply the URL filtering profile that you just enabled for Safe Search Enforcement.<br>2. On the **Actions** tab, select the **URL Filtering** profile.<br>3. Click **OK** to save the security policy rule. |


© Palo Alto Networks, Inc.

| Enable Transparent Safe Search Enforcement (Continued) | | |
|---|---|---|
| Step 3 | (Optional, but recommended) Block Bing search traffic running over SSL.<br><br>Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL. | 1. Add a custom URL category for Bing:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as EnableBingSafeSearch.<br>c. **Add** the following to the Sites list:<br>`www.bing.com/images/*`<br>`www.bing.com/videos/*`<br>d. Click **OK** to save the custom URL category object.<br>2. Create another URL filtering profile to block the custom category you just created:<br>a. Select **Objects > Security Profiles > URL Filtering**.<br>b. **Add** a new profile and give it a descriptive **Name**.<br>c. Locate the custom category you just created in the Category list and set it to **block**.<br>d. Click **OK** to save the URL filtering profile.<br>3. **Add** a security policy rule to block Bing SSL traffic:<br>a. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.<br>b. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.<br>c. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.<br>d. Select **TCP** as the **Protocol**, set the **Destination Port** to `443`.<br>e. Click **OK** to save the rule.<br>f. Use the **Move** options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled. |
| Step 4 | Edit the URL Filtering Safe Search Block Page, replacing the existing code with the Javascript for rewriting search query URLs to enforce safe search transparently.<br><br> For a script that you can copy and paste, go here. | 1. Select **Device > Response Pages > URL Filtering Safe Search Block Page**.<br>2. Select **Predefined** and then click **Export** to save the file locally.<br>3. Use an HTML editor and replace all of the existing block page text with the following text and then save the file: |

**Enable Transparent Safe Search Enforcement (Continued)**

```
<html>
<head>
<title>Search Blocked</title>
<meta http-equiv="Content-Type" content="text/html; charset=utf-8">
<META HTTP-EQUIV="PRAGMA" CONTENT="NO-CACHE">
<meta name="viewport" content="initial-scale=1.0">
<style>
  #content {
    border:3px solid#aaa;
    background-color:#fff;
    margin:1.5em;
    padding:1.5em;
    font-family:Tahoma,Helvetica,Arial,sans-serif;
    font-size:1em;
  }
  h1 {
    font-size:1.3em;
    font-weight:bold;
    color:#196390;
  }
  b {
    font-weight:normal;
    color:#196390;
  }
</style>
</head>
    <body bgcolor="#e7e8e9">
    <h1>Search Blocked</h1>
        <p><b>User:</b> <user/> </p>
       <p>Your search results have been blocked because your search settings are not in accordance with company
policy.  In order to continue, please update your search settings so that Safe Search is set to the strictest
setting. If you are currently logged into your account, please also lock Safe Search and try your search
again.</p>
        <p>For more information, please refer to: <a href="<ssurl/>"><ssurl/></a></p>
        <p id="java_off"> Please enable JavaScript in your browser.<br></p>
        <p><b>Please contact your system administrator if you believe this message is in error.</b></p>


</body>
<script>

    // Grab the URL that's in the browser.
    var s_u = location.href;

    //bing
    // Matches the forward slashes in the beginning, anything, then ".bing." then anything followed by a non
greedy slash. Hopefully the first forward slash.
    var b_a = /^.*\/\/(.+\.bing\..+?)\//.exec(s_u);

    if (b_a) {

        s_u = s_u + "&adlt=strict";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You are being redirected to a safer search!';

    }

    //google
    // Matches the forward slashes in the beginning, anything, then ".google." then anything followed by a non
greedy slash. Hopefully the first forward slash.
    var g_a = /^.*\/\/(.+\.google\..+?)\//.exec(s_u);

    if (g_a) {

     s_u = s_u.replace(/&safe=off/ig,"");
        s_u = s_u + "&safe=active";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You are being redirected to a safer search!';

    }

    //yahoo
    // Matches the forward slashes in the beginning, anything, then ".yahoo."" then anything followed by a non
greedy slash. Hopefully the first forward slash.
    var y_a = /^.*\/\/(.+\.yahoo\..+?)\//.exec(s_u);
    if (y_a) {

        s_u = s_u.replace(/&vm=p/ig,"");
        s_u = s_u + "&vm=r";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You are being redirected to a safer search!';
    }
    document.getElementById("java_off").innerHTML = ' ';

    </script>
</html>
```

 © Palo Alto Networks, Inc.

| Enable Transparent Safe Search Enforcement (Continued) | | |
|---|---|---|
| Step 5 | Import the edited URL Filtering Safe Search Block page onto the firewall. | 1. To import the edited block page, select **Device > Response Pages > URL Filtering Safe Search Block Page**.<br>2. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.<br>3. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select **shared** to make it available to all virtual systems.<br>4. Click **OK** to import the file. |
| Step 6 | Enable SSL Forward Proxy decryption.<br><br>Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings. | 1. Add a custom URL category for the search sites:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as SearchEngineDecryption.<br>c. **Add** the following to the Sites list:<br>`www.bing.*`<br>`www.google.*`<br>`search.yahoo.*`<br>d. Click **OK** to save the custom URL category object.<br>2. Follow the steps to Configure SSL Forward Proxy.<br>3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**. |
| Step 7 | Save the configuration. | Click **Commit**. |

# URL Filtering Use Case Examples

The following use cases show how to use App-ID to control a specific set of web-based applications and how to use URL categories as match criteria in a policy. When working with App-ID, it is important to understand that each App-ID signature may have dependencies that are required to fully control an application. For example, with Facebook applications, the App-ID facebook-base is required to access the Facebook website and to control other Facebook applications. For example, to configure the firewall to control Facebook email, you would have to allow the App-IDs facebook-base and facebook-mail. As another example, if you search Applipedia (the App-ID database) for LinkedIn, you will see that in order to control LinkedIn mail, you need to apply the same action to both App-IDs: linkedin-base and linkedin-mail. To determine application dependencies for App-ID signatures, visit Applipedia, search for the given application, and then click the application for details.

The User-ID feature is required to implement policies based on users and groups and a Decryption policy is required to identify and control websites that are encrypted using SSL.

This section includes two uses cases:

- Use Case: Control Web Access
- Use Case: Use URL Categories in Policy

 © Palo Alto Networks, Inc.

## Use Case: Control Web Access

When using URL filtering to control user website access, there may be instances where granular control is required for a given website. In this use case, a URL filtering policy is applied to the security policy that allows web access for your users and the *social-networking* URL category is set to block, but the allow list in the URL profile is configured to allow the social networking site Facebook. To further control Facebook, the company policy also states that only marketing has full access to Facebook and all other users within the company can only read Facebook posts and cannot use any other Facebook applications, such as email, posting, chat, and file sharing. To accomplish this requirement, App-ID must be used to provide granular control over Facebook.

The first security rule will allow marketing to access the Facebook website as well as all Facebook applications. Because this allow rule will also allow access to the Internet, threat prevention profiles are applied to the rule, so traffic that matches the policy will be scanned for threats. This is important because the allow rule is terminal and will not continue to check other rules if there is a traffic match.

| **Control Web Access** | | |
|---|---|---|
| Step 1 | Confirm that URL filtering is licensed. | 1. Select **Device > Licenses** and confirm that a valid date appears for the URL filtering database that will used. This will either be PAN-DB or BrightCloud.<br>2. If a valid license is not installed, see Enable PAN-DB URL Filtering. |
| Step 2 | Confirm that User-ID is working. User-ID is required to create policies based on users and groups. | 1. To check group mapping, from the CLI, enter the following command:<br>`show user group-mapping statistics`<br>2. To check user mapping, from the CLI, enter the following command:<br>`show user ip-user-mapping-mp all`<br>3. If statistics do not appear and/or IP to user mapping information is not displayed, see User-ID. |
| Step 3 | Set up a URL filtering profile by cloning the default profile. | 1. Select **Objects > Security Profiles > URL Filtering** and select the **default** profile.<br>2. Click the **Clone** icon. A new profile should appear named **default-1**.<br>3. Select the new profile and rename it. |

| Control Web Access (Continued) | | |
|---|---|---|
| Step 4 | Configure the URL filtering profile to block social-networking and allow Facebook. | 1. Modify the new URL filtering profile and in the **Category** list scroll to **social-networking** and in the **Action** column click on **allow** and change the action to **block**.<br>2. In the **Allow List** box, type `facebook.com`, press enter to start a new line and then type `*.facebook.com`. Both of these formats are required, so all URL variants a user may use will be identified, such as facebook.com, www.facebook.com, and https://facebook.com.<br> 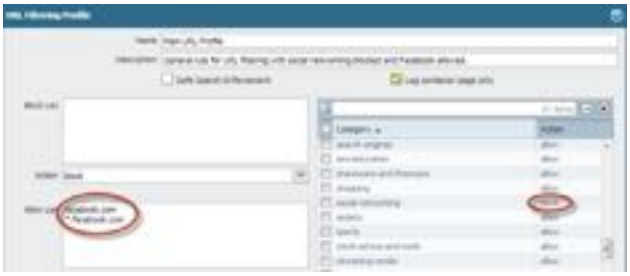 <br>3. Click **OK** to save the profile. |
| Step 5 | Apply the new URL filtering profile to the security policy rule that allows web access from the user network to the Internet. | 1. Select **Policies > Security** and click on the policy rule that allows web access.<br>2. On the **Actions** tab, select the URL profile you just created from the **URL Filtering** drop-down.<br> 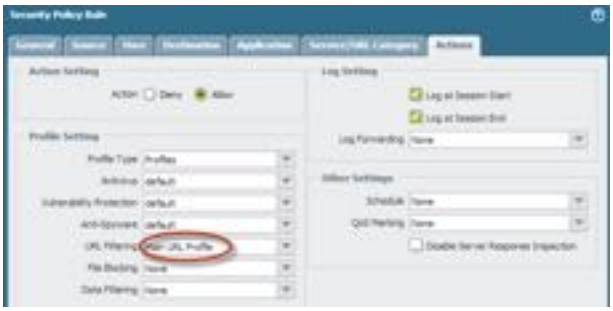 <br>3. Click **OK** to save. |


© Palo Alto Networks, Inc.

**Control Web Access (Continued)**

Step 6 Create the security policy that will allow marketing access the Facebook website and all Facebook applications.

This rule must precede other rules because it is more specific than the other policies and because it is an allow rule, which will terminate when a traffic match occurs.

1. Select **Policies > Security** and click **Add**.
2. Enter a **Name** and optionally a **Description** and **Tag**(s).
3. On the **Source** tab add the zone where the users are connected.
4. On the **User** tab in the **Source User** section click **Add**.
5. Select the directory group that contains your *marketing* users.
6. On the **Destination** tab, select the zone that is connected to the Internet.
7. On the **Applications** tab, click **Add** and add the *facebook* App-ID signature.
8. On the **Actions** tab, add the default profiles for **Antivirus**, **Vulnerability Protection**, and **Anti-Spyware**.

9. Click **OK** to save the security profile.

   The *facebook* App-ID signature used in this policy encompasses all Facebook applications, such as facebook-base, facebook-chat, and facebook-mail, so this is the only App-ID signature required in this rule.

   With this policy in place, when a marketing employee attempts to access the Facebook website or any Facebook application, the rule matches based on the user being part of the marketing group. For traffic from any user outside of marketing, the rule will be skipped because there would not be a traffic match and rule processing would continue.

| Control Web Access (Continued) | |
|---|---|
| Step 7 Configure the security policy to block all other users from using any Facebook applications other than simple web browsing. The easiest way to do this is to clone the marketing allow policy and then modify it. | 1. From **Policies > Security** click the marketing Facebook allow policy you created earlier to highlight it and then click the **Clone** icon.<br>2. Enter a **Name** and optionally enter a **Description** and **Tag**('s).<br>3. On the **User** tab highlight the marketing group and delete it and in the drop-down select **any**.<br>4. On the **Applications** tab, click the *facebook* App-ID signature and delete it.<br>5. Click **Add** and add the following App-ID signatures:<br>• facebook-apps<br>• facebook-chat<br>• facebook-file-sharing<br>• facebook-mail<br>• facebook-posting<br>• facebook-social-plugin<br>6. On the **Actions** tab in the **Action Setting** section, select **Deny**. The profile settings should already be correct because this rule was cloned. |
| | <br>7. Click **OK** to save the security profile.<br>8. Ensure that this new deny rule is listed after the marketing allow rule, to ensure that rule processing occurs in the correct order to allow marketing users and then to deny/limit all other users.<br>9. Click **Commit** to save the configuration. |

With these policies in place, any user who is part of the marketing group will have full access to all Facebook applications and any user that is not part of the marketing group will only have read-only access to the Facebook website and will not be able to use Facebook functions such as post, chat, email, and file sharing.

 © Palo Alto Networks, Inc.

## Use Case: Use URL Categories in Policy

URL categories can also be used as match criteria in the following policy types: captive portal, decryption, security, and QoS. In this use case, URL categories will be used in decryption policies to control which web categories should be decrypted or not decrypted. The first rule is a no-decrypt rule that will not decrypt user traffic if the website category is *financial-services* or *health-and-medicine* and the second rule will decrypt all other traffic. The decryption policy type is *ssl-forward-proxy*, which is used for controlling decryption for all outbound connections performed by users.

**Configure a Decryption Policy Based on URL Category**

Step 1 Create the no-decrypt rule that will be listed first in the decryption policies list. This will prevent any website that is in the *financial-services* or *health-and-medicine* URL categories from being decrypted.

1. Select **Policies > Decryption** and click **Add**.
2. Enter a **Name** and optionally enter a **Description** and **Tag**(s).
3. On the **Source** tab, add the zone where the users are connected.
4. On the **Destination** tab, enter the zone that is connected to the Internet.
5. On the **URL Category** tab, click **Add** and select the *financial-services* and *health-and-medicine* URL categories.
6. On the **Options** tab, set the action to **No Decrypt** and the **Type** to **SSL Forward Proxy**.

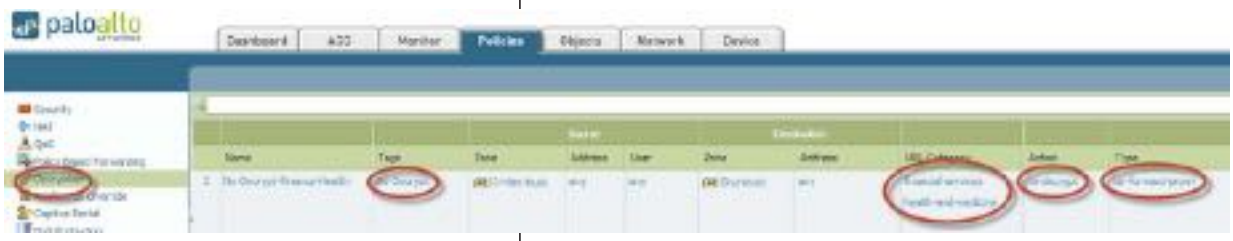


7. Click **OK** to save the policy.

| Configure a Decryption Policy Based on URL Category (Continued) | |
|---|---|
| Step 2 Create the decryption policy that will decrypt all other traffic. This policy will be listed after the no-decrypt policy. | 1. Select the no-decrypt policy you created previously and then click **Clone**.<br>2. Enter a **Name** and optionally enter a **Description** and **Tag**(s).<br>3. On the **URL Category** tab, select *financial-services* and *health-and-medicine* and then click the **Delete** icon.<br>4. On the **Options** tab, set the action to **Decrypt** and the **Type** to **SSL Forward Proxy**. |
| | 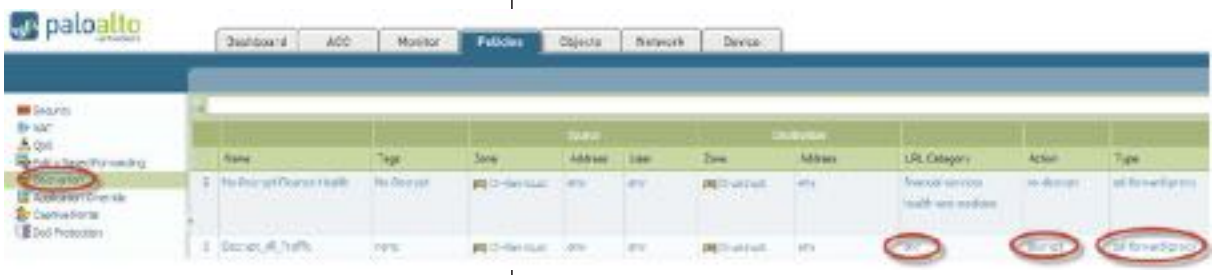 <br>5. Ensure that this new decryption rule is listed after the no-decrypt rule as shown in the previous screen capture. This will ensure that rule processing occurs in the correct order, so websites in the *financial-services* and *health-and-medicine* are not decrypted<br>6. Click **OK** to save the policy. |
| Step 3 (BrightCloud only) Enable cloud lookups for dynamically categorizing a URL when the category is not available on the local database on the firewall. | 1. Access the CLI on the firewall.<br>2. Enter the following commands to enable Dynamic URL Filtering:<br>a. **configure**<br>b. **set deviceconfig setting url dynamic-url yes**<br>c. commit |
| Step 4 Save the configuration. | Click **Commit**. |

With these two decrypt policies in place, any traffic destined for the *financial-services* or *health-and-medicine* URL categories will not be decrypted. All other traffic will be decrypted.

You can also define more granular control over decryption policies by defining decryption profiles, which are used to perform checks such as server certificate checks or blocking sessions with expired certificates. The profile is then added in the **Options** tab of the decryption policy. For a complete list of checks that can be performed, select **Objects > Decryption Profiles** from the firewall and then click the help icon.

Now that you have a basic understanding of the powerful features of URL filtering, App-ID, and User-ID, you can apply similar policies to your firewall to control any application in the Palo Alto Networks App-ID signature database and control any website contained in the URL filtering database.

For help in troubleshooting URL filtering issues, see Troubleshoot URL Filtering.

 © Palo Alto Networks, Inc.

# Troubleshoot URL Filtering

The following topics provide troubleshooting guidelines for diagnosing and resolving common URL filtering problems.


- Problems Activating PAN-DB
- PAN-DB Cloud Connectivity Issues
- URLs Classified as Not-Resolved
- Incorrect Categorization
- URL Database Out of Date


## Problems Activating PAN-DB

The following table describes procedures that you can use to resolve issues with activating PAN-DB.

| Troubleshoot PAN-DB Activation Issues |
|---|
| 1. Access the CLI on the firewall. |
| 2. Verify whether PAN-DB has been activated by running the following command:<br>`admin@PA-200> show system setting url-database`<br>If the response is `paloaltonetworks`, then PAN-DB is the active vendor. |
| 3. Verify that the firewall has a valid PAN-DB license by running the following command:<br>`admin@PA-200> request license info`<br>You should see the license entry Feature: PAN_DB URL Filtering. If the license is not installed, you will need to obtain and install a license. See Configure URL Filtering. |
| 4. After the license is installed, download a new PAN-DB seed database by running the following command:<br>`admin@PA-200> request url-filtering download paloaltonetworks region <region>` |
| 5. Check the download status by running the following command:<br>`admin@PA-200> request url-filtering download status vendor paloaltonetworks`<br>• If the message is different from `PAN-DB download: Finished successfully`, stop here; there may be a problem connecting to the cloud. Attempt to solve the connectivity issue by performing basic network troubleshooting between the firewall and the Internet. For more information, see PAN-DB Cloud Connectivity Issues.<br>• If the message is `PAN-DB download: Finished successfully`, the firewall successfully downloaded the URL seed database. Try to enable PAN-DB again by running the following command:<br>`admin@PA-200> set system setting url-database paloaltonetworks` |
| 6. If the problems persists, contact Palo Alto Networks support. |

## PAN-DB Cloud Connectivity Issues

To check cloud connectivity, run the following command:

```
admin@pa-200> show url-cloud status
```

If the cloud is accessible, the expected response is similar to the following:

```
admin@PA-200> show url-cloud status
PAN-DB URL Filtering
License :                          valid
Current cloud server :             s0000.urlcloud.paloaltonetworks.com
Cloud connection :                 connected
URL database version - device :    2013.11.18.000
URL database version - cloud :     2013.11.18.000  ( last update time
2013/11/19
 13:20:51 )
URL database status :              good
URL protocol version - device :    pan/0.0.2
URL protocol version - cloud :     pan/0.0.2
Protocol compatibility status :    compatible
```

If the cloud is note accessible, the expected response is similar to the following:

```
admin@PA-200> show url-cloud status
PAN-DB URL Filtering
License :                          valid
Cloud connection :                 not connected
URL database version - device :    2013.11.18.000
URL database version - cloud :     2013.11.18.000  ( last update time
2013/11/19
 13:20:51 )
URL database status :              good
URL protocol version - device :    pan/0.0.2
URL protocol version - cloud :     pan/0.0.2
Protocol compatibility status :    compatible
```

 © Palo Alto Networks, Inc.

The following table describes procedures that you can use to resolve issues based on the output of the show url-cloud status command, how to ping the URL cloud servers, and what to check if the firewall is in a High Availability (HA) configuration.

**Troubleshoot Cloud Connectivity Issues**

- PAN-DB URL Filtering license field shows invalid—Obtain and install a valid PAN-DB license.
- URL database status is out of date—Download a new seed database by running the following command:

  ```
  admin@pa-200> request url-filtering download paloaltonetworks region <region>
  ```
- URL protocol version shows not compatible—Upgrade PAN-OS to the latest version.
- Attempt to ping the PAN-DB cloud server from the firewall by running the following command:

  ```
  admin@pa-200> ping source <ip-address> host s0000.urlcloud.paloaltonetworks.com
  ```
  For example, if your management interface IP address is 10.1.1.5, run the following command:
  ```
  admin@pa-200> ping source 10.1.1.5 host s0000.urlcloud.paloaltonetworks.com
  ```
- If the firewall is in an HA configuration, verify that the HA state of the devices supports connectivity to the cloud systems. You can determine the HA state by running the following command:

  ```
  admin@pa-200> show high-availability state
  ```
  Connection to the cloud will be blocked if the firewall is not in one of the following states:
  - active
  - active-primary
  - active-secondary

If the problem persists, contact Palo Alto Networks support.

## URLs Classified as Not-Resolved

The following table describes procedures you can use to resolve issues where some or all of the URLs being identified by PAN-DB are classified as *Not-resolved*:

**Troubleshoot URLs Classified as Not-Resolved**

1. Check the PAN-DB cloud connection by running the following command:

   ```
   admin@PA-200> show url-cloud status
   ```
   The Cloud connection: field should show *connected*. If you see anything other than *connected*, any URL that do not exist in the management plane cache will be categorized as *not-resolved*. To resolve this issue, see PAN-DB Cloud Connectivity Issues.
2. If the cloud connection status shows *connected*, check the current utilization of the firewall. If the firewall's performance is spiking, URL requests may be dropped (may not reach the management plane), and will be categorized as *not-resolved*.

   To view system resources, run the following command and view the *%CPU* and *%MEM* columns:

   ```
   admin@PA-200> show system resources
   ```
   You can also view system resources from the firewall's web interfaces by clicking the **Dashboard** tab and viewing the **System Resources** section.
3. If the problem persist, contact Palo Alto Networks support.

## Incorrect Categorization

The following steps describe the procedures you can use if you identify a URL that does not have the correct categorization. For example, if the URL *paloaltonetworks.com* was categorized as *alcohol-and-tobacco*, the categorization is not correct; the category should be *computer-and-internet-info*.

**Troubleshoot Incorrect Categorization Issues**

1. Verify the category in the dataplane by running the following command:

   `admin@PA-200>` show running url <*URL*>

   For example, to view the category for the Palo Alto Networks website, run the following command:

   ```
   admin@PA-200> show running url paloaltonetworks.com
   ```

   If the URL stored in the dataplane cache has the correct category (computer-and-internet-info in this example), then the categorization is correct and no further action is required. If the category is not correct, continue to the next step.

2. Verify if the category in the management plane by running the command:

   `admin@PA-200> test url-info-host` <*URL*>

   For example:

   ```
   admin@PA-200> test url-info-host paloaltonetworks.com
   ```

   If the URL stored in the management plane cache has the correct category, remove the URL from the dataplane cache by running the following command:

   `admin@PA-200> clear url-cache url` <*URL*>

   The next time the device requests the category for this URL, the request will be forwarded to the management plane. This will resolve the issue and no further action is required. If this does not solve the issue, go to the next step to check the URL category on the cloud systems.

3. Verify the category in the cloud by running the following command:

   `admin@PA-200> test url-info-cloud <URL>`

4. If the URL stored in the cloud has the correct category, remove the URL from the dataplane and the management plane caches.

   Run the following command to delete a URL from the dataplane cache:

   ```
   admin@PA-200> clear url-cache url <URL>
   ```

   Run the following command to delete a URL from the management plane cache:

   ```
   admin@PA-200> delete url-database url <URL>
   ```

   The next time the device queries for the category of the given URL, the request will be forwarded to the management plane and then to the cloud. This should resolve the category lookup issue. If problems persist, see the next step to submit a categorization change request.

5. To submit a change request from the web interface, go to the URL log and select the log entry for the URL you would like to have changed.

 © Palo Alto Networks, Inc.

**Troubleshoot Incorrect Categorization Issues**

6. Click the **Request Categorization** change link and follow instructions. You can also request a category change from the Palo Alto Networks Test A Site website by searching for the URL and then clicking the **Request Change** icon. To view a list of all available categories with descriptions of each category, refer to https://urlfiltering.paloaltonetworks.com/CategoryList.aspx.

   If your change request is approved, you will receive an email notification. You then have two options to ensure that the URL category is updated on the firewall:

   - Wait until the URL in the cache expires and the next time the URL is accessed by a user, the new categorization update will be put in the cache.
   - Run the following command to force an update in the cache:

     ```
     admin@PA-200> request url-filtering update url <URL>
     ```

## URL Database Out of Date

If you have observed through the syslog or the CLI that PAN-DB is out-of-date, it means that the connection from the firewall to the URL Cloud is blocked. This usually occurs when the URL database on the firewall is too old (version difference is more than three months) and the cloud cannot update the firewall automatically. In order to resolve this issue, you will need to re-download an initial seed database from the cloud (this operation is not blocked). This will result in an automatic re-activation of PAN-DB.

To manually update the database, perform one of the following steps:

- From the web interface, select **Device > Licenses** and in the **PAN-DB URL Filtering** section click the **Re-Download** link.
- From the CLI, run the following command:

  ```
  admin@PA-200> request url-filtering download paloaltonetworks region <region_name>
  ```

When the seed database is re-download, the URL cache in the management plane and dataplane will be purged. The management plane cache will then be re-populated with the contents of the new seed database.

 © Palo Alto Networks, Inc.

# Quality of Service

Quality of Service (QoS) is a set of technologies that work on a network to guarantee its ability to dependably run high-priority applications and traffic under limited network capacity. QoS technologies accomplish this by providing differentiated handling and capacity allocation to specific flows in network traffic. This enables the network administrator to assign the order in which traffic is handled, and the amount of bandwidth afforded to traffic.

Palo Alto Networks Application Quality of Service (QoS) provides basic QoS applied to networks and extends it to provide QoS to applications and users.

Use the following topics to learn about and configure Palo Alto Networks Application QoS:

- QoS Overview
- QoS Concepts
- Configure QoS
- Configure QoS for a Virtual System
- QoS Use Case Examples

- You can use the Palo Alto Networks product comparison tool to view the QoS features supported on your firewall platform. Select two or more product platforms and click **Compare Now** to view QoS feature support for each platform (for example, you can check if your firewall platform supports QoS on subinterfaces and if so, the maximum number of subinterfaces that QoS can be enabled on).
- The PA-7050 firewall supports QoS on Aggregate Ethernet (AE) interfaces as of PAN-OS 6.0.0.

# QoS Overview

Use QoS to prioritize and adjust quality aspects of network traffic. You can assign the order in which packets are handled and allot bandwidth, ensuring preferred treatment and optimal levels of performance are afforded to selected traffic, applications, and users.

Service quality measurements subject to a QoS implementation are bandwidth (maximum rate of transfer), throughput (actual rate of transfer), latency (delay), and jitter (variance in latency). The capability to shape and control these service quality measurements makes QoS of particular importance to high-bandwidth, real-time traffic such as voice over IP (VoIP), video conferencing, and video-on-demand that has a high sensitivity to latency and jitter. Additionally, use QoS to achieve outcomes such as the following:

- Prioritize network and application traffic, guaranteeing high priority to important traffic or limiting non-essential traffic.
- Achieve equal bandwidth sharing among different subnets, classes, or users in a network.
- Allocate bandwidth externally or internally or both, applying QoS to both upload and download traffic or to only upload or download traffic.
- Ensure low latency for customer and revenue-generating traffic in an enterprise environment.
- Perform traffic profiling of applications to ensure bandwidth usage.

QoS implementation on a Palo Alto Networks firewall begins with three primary configuration components that support a full QoS solution: a QoS Profile, a QoS Policy, and setting up the QoS Egress Interface. Each of these options in the QoS configuration task facilitate a broader process that optimizes and prioritizes the traffic flow and allocates and ensures bandwidth according to configurable parameters.

The figure Figure: QoS Traffic Flow shows traffic as it flows from the source, is shaped by the firewall with QoS enabled, and is ultimately prioritized and delivered to its destination.

**Figure: QoS Traffic Flow**

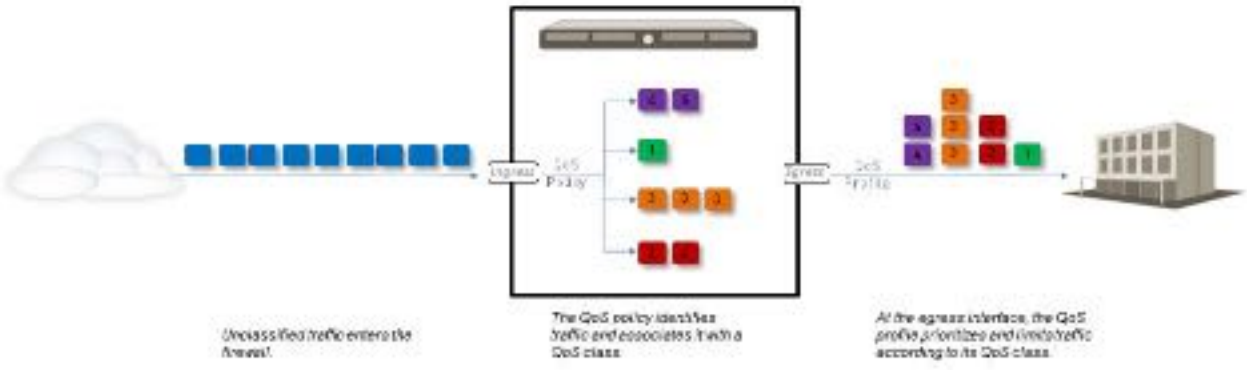

The QoS configuration options allow you to control the traffic flow and define it at different points in the flow. The Figure: QoS Traffic Flow indicates where the configurable options define the traffic flow. Use the QoS Profile to define QoS classes and use the QoS Policy to associate QoS classes with selected traffic. Enable the QoS Profile on an interface to shape traffic according to the QoS configuration as it flows through the network.

 © Palo Alto Networks, Inc.

You can configure a QoS Profile and QoS Policy individually or in any order, according to your preference. Each of the QoS configuration options has components that influence the definition of the other options and the QoS configuration options can be used to create a full and granular QoS policy or can be used sparingly with minimal administrator action.

Each firewall model supports a maximum number of ports that can be configured with QoS. Refer to the spec sheet for your firewall model or use the product comparison tool to view QoS feature support for two or more firewalls on a single page.

# QoS Concepts

Use the following topics to learn about the different components and mechanisms of a QoS configuration on a Palo Alto Networks firewall:


- QoS for Applications and Users
- QoS Profile
- QoS Classes
- QoS Policy
- QoS Egress Interface
- QoS Cleartext and Tunneled Traffic


## QoS for Applications and Users

A Palo Alto Networks firewall provides basic QoS, controlling traffic leaving the firewall according to network or subnet, and extends the power of QoS to also classify and shape traffic according to application and user. The Palo Alto Networks firewall provides this capability by integrating the features App-ID and User-ID with the QoS configuration. App-ID and User-ID entries that exist to identify specific applications and users in your network are available in the QoS configuration so that you can easily specify applications and users to apply QoS to.

You can use a QoS Policy in the web interface (**Policies** > **QoS**) to apply QoS specifically to an application's traffic:

Or to a user's traffic:

See App-ID and User-ID for more information on these features.


© Palo Alto Networks, Inc.

## QoS Profile

Use a QoS profile to define values of up to eight QoS classes contained within that single profile (**Network** > **Network Profiles** > **QoS Profile**):

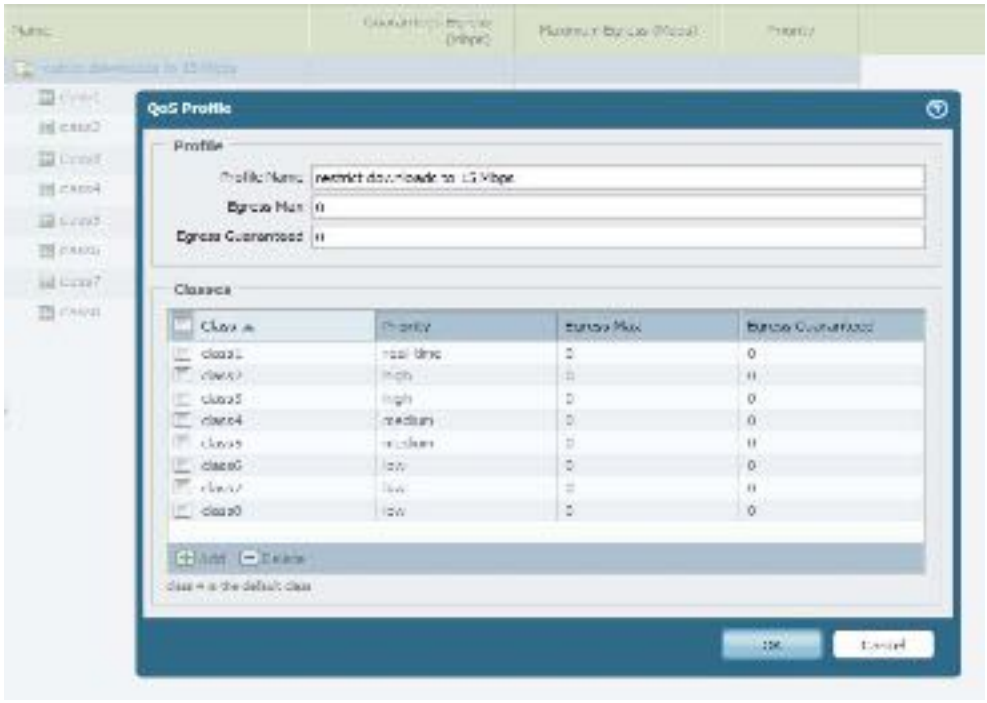


You enable QoS by applying a QoS profile to the egress interface for network, application or user traffic (or specifically, for cleartext or tunneled traffic). An interface configured with QoS shapes traffic according to the QoS profile class definitions and the traffic associated with those classes in the QoS policy.

A default QoS Profile is available on the firewall. The default profile and the classes defined in the profile do not have predefined maximum or guaranteed bandwidth limits.

You can set bandwidth limits for a QoS profile and/or set limits for individual QoS classes within the QoS profile. The total guaranteed bandwidth limits of all eight QoS classes in a QoS Profile cannot exceed the total bandwidth allocated to that QoS Profile. Enabling QoS on a physical interface includes setting the maximum bandwidth for traffic leaving the firewall through this interface. A QoS profile's guaranteed bandwidth (the **Egress Guaranteed** field) should not exceed the bandwidth allocated to the physical interface that QoS is enabled on.

For details, see Create a QoS profile.

## QoS Classes

A QoS class determines the priority and bandwidth for traffic it is assigned to. In the web interface, use the QoS profile to define QoS classes (**Network** > **Network Profiles** > **QoS Profile**):

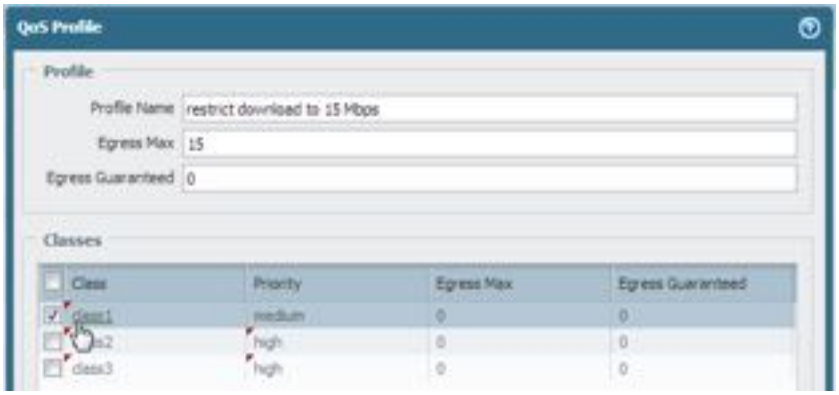


Defining a QoS class includes setting the class's Priority, maximum bandwidth (Egress Max), and guaranteed bandwidth (Egress Guaranteed).

Real-time priority is typically used for applications that are particularly sensitive to latency, such as voice and video applications.

Use the QoS policy to assign a QoS class to specified traffic (**Policies** > **QoS**):

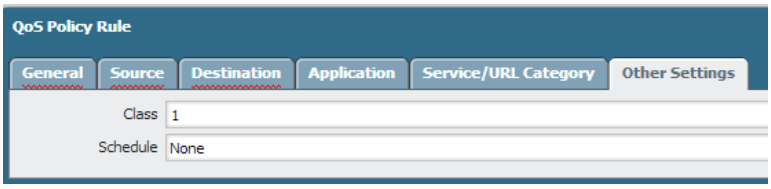


There are up to eight definable QoS classes in a single QoS profile. Unless otherwise configured, traffic that does not match a QoS class is assigned a class of 4.

QoS priority queuing and bandwidth management, the fundamental mechanisms of a QoS configuration, are configured within the QoS class definition (see Step 3). Queuing priority is determined by the priority set for a QoS class. Bandwidth management is determined according to the maximum and guaranteed bandwidths set for a QoS class.

 © Palo Alto Networks, Inc.

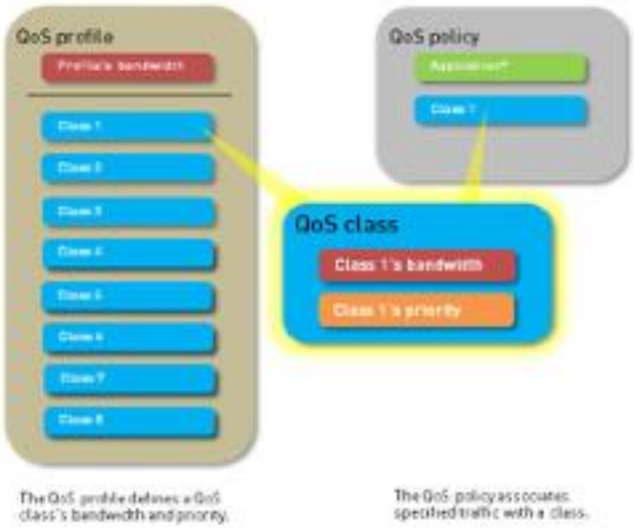


The queuing and bandwidth management mechanisms determine the order of traffic and how traffic is handled upon entering or leaving a network:

- **QoS Priority:** One of four QoS priorities can be defined in a QoS class: real-time, high, medium, and low. When a QoS class is associated with specific traffic, the priority defined in that QoS class is assigned to the traffic. Packets in the traffic flow are then queued according to their priority until the network is ready to process them. This method of priority queuing provides the capability to ensure that important traffic, applications, or users takes precedence.
- **QoS Class Bandwidth Management:** QoS class bandwidth management provides the capability to control traffic flows on a network so that traffic does not exceed network capacity, resulting in network congestion, or to allocate specific bandwidth limits to traffic, applications, or users. You can set overall limits on bandwidth using the QoS profile or set limits for individual QoS classes. A QoS profile and QoS classes in the profile have guaranteed and maximum bandwidth limits. The guaranteed bandwidth limit (Egress Guaranteed) ensures that any amount of traffic up to that set bandwidth limit is processed. The maximum bandwidth limit (Egress Max) sets the total limit of bandwidth allocated to either the QoS Profile or QoS Class. Traffic in excess of the Maximum Bandwidth limit is dropped. The total bandwidth limits and guaranteed bandwidth limits of QoS classes in a QoS profile cannot exceed the bandwidth limit of the QoS profile.

## QoS Policy

In a QoS configuration, the QoS policy identifies traffic requiring QoS treatment (either preferential treatment or bandwidth-limiting) using a defined parameter or multiple parameters and assigns it a class.

Use the QoS Policy, similar to a security policy, to set the criteria that identifies traffic:

- Applications and application groups.
- Source zones, source addresses, and source users.
- Destination zones and destination addresses.

- Services and service groups limited to specific TCP and/or UDP port numbers.
- URL categories, including custom URL categories.

The QoS Policy in the web interface (**Policies** > **QoS**) allows you to associate the criteria used to specify traffic with a QoS class.

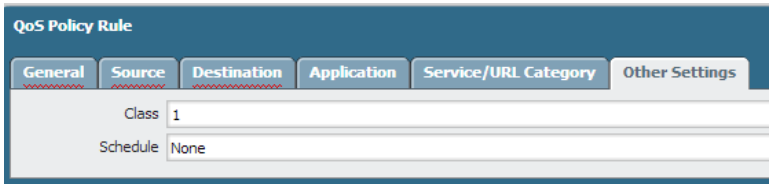


## QoS Egress Interface

Enabling a QoS profile on the egress interface of the traffic identified for QoS treatment completes a QoS configuration. The ingress interface for QoS traffic is the interface on which the traffic enters the firewall. The egress interface for QoS traffic is the interface that traffic leaves the firewall from. QoS is always enabled and enforced on a traffic flow's egress interface. The egress interface in a QoS configuration can either be the external- or internal-facing interface of the firewall, depending on the flow of the traffic receiving QoS treatment.

For example, in an enterprise network, if you are limiting employees' download traffic from a specific website, the egress interface in the QoS configuration is the firewall's internal interface, as the traffic flow is from the Internet, through the firewall, and to your company network. Alternatively, when limiting employees' upload traffic to the same website, the egress interface in the QoS configuration is the firewall's external interface, as the traffic you are limiting flows from your company network, through the firewall, and then to the Internet.

See Step 3 to learn how to Identify the egress interface for applications that you identified as needing QoS treatment.

 © Palo Alto Networks, Inc.

## QoS Cleartext and Tunneled Traffic

Within the QoS physical interface configuration, you can provide more granular QoS settings for cleartext traffic and tunneled traffic leaving through the interface. Individual tunnel interfaces can be assigned different QoS Profiles. Cleartext traffic can be assigned different QoS Profiles according to traffic's source interface and source subnet. In this case, a source interface and source subnet can be associated with a QoS Profile. If you choose not to select cleartext or tunneled traffic for unique QoS treatment, enabling QoS on an interface requires selecting a default QoS Profile to determine how to shape traffic for specific tunnel interfaces or, in the case of cleartext traffic, source interfaces and source subnets.

On Palo Alto Networks firewalls, the term "tunneled traffic" refers to tunnel interface traffic, specifically IPSec traffic in tunnel mode.

# Configure QoS

Use the following task to configure Quality of Service (QoS), including how to create a QoS profile, create a QoS policy, and enable QoS on an interface.

| Configure QoS | | |
|---|---|---|
| Step 1 | Identify the traffic to which to apply QoS.<br><br>This example shows how to use QoS to limit web browsing. | Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.<br><br>Click any application name to display detailed application information. |
| Step 2 | Identify the egress interface for applications that you identified as needing QoS treatment.<br><br>**Tip:** The egress interface for traffic depends on the traffic flow. If you are shaping incoming traffic, the egress interface is the internal-facing interface. If you are shaping outgoing traffic, the egress interface is the external-facing interface. | Select **Monitor > Logs > Traffic** to view the device's traffic logs.<br><br>To filter and only show logs for a specific application:<br>• If an entry is displayed for the application, click the underlined link in the Application column then click the Submit icon.<br>• If an entry is not displayed for the application, click the Add Log icon and search for the application.<br><br>The **Egress I/F** in the traffic logs displays each application's egress interface. To display the **Egress I/F** column if it is not displayed by default:<br>• Click any column header to add a column to the log:<br> 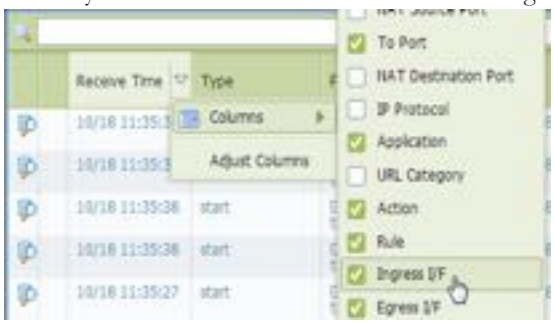 <br><br>• Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface listed in the Destination section:<br> <br><br>In this example, the egress interface for web-browsing traffic is ethernet 1/1. |

 © Palo Alto Networks, Inc.

**Configure QoS (Continued)**

Step 3 Create a QoS profile.

You can edit any existing QoS profile, including the default, by clicking the QoS profile name.

1. Select **Network > Network Profiles > QoS Profile** and click **Add** to open the **QoS Profile** dialog.
2. Enter a descriptive **Profile Name**.
3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS Profile.
4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS Profile.

   Any traffic that exceeds the QoS Profile's egress guaranteed limit is best effort but is not guaranteed.
5. In the Classes section, specify how to treat up to eight individual QoS classes:
   a. Click **Add** to add a class to the QoS Profile.
   b. Select the **Priority** for the class.
   c. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.
   d. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.
6. Click **OK** to save the QoS Profile.

In the following example, the QoS Profile named Limit Web Browsing limits traffic identified as Class 2 traffic to maximum bandwidth of 50 Mbps and a guaranteed bandwidth of 2 Mbps. Any traffic that is associated with class 2 in a the QoS policy (Step 4) is subject to these limits.

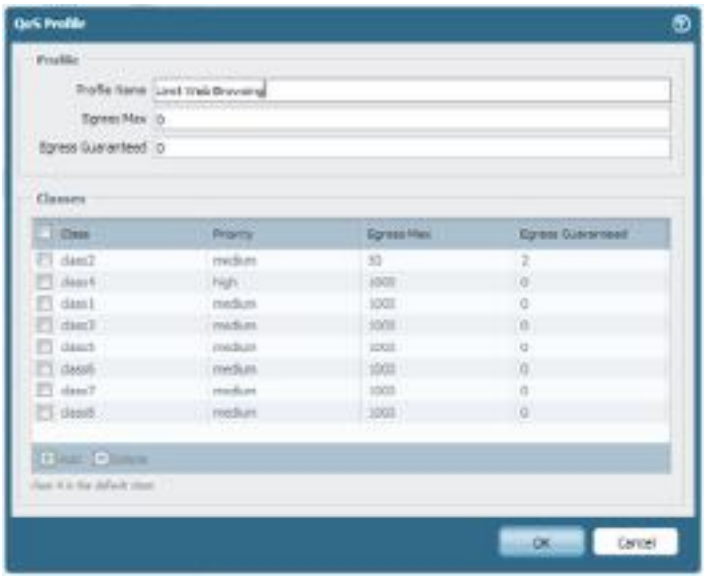

| Configure QoS (Continued) | |
|---|---|
| Step 4 Create a QoS policy. | 1. Select **Policies > QoS** and click **Add** to open the QoS Policy Rule dialog.<br>2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**.<br>3. Specify the traffic to which the QoS Policy Rule will apply. Use the **Source**, **Destination**, **Application**, and **Service/URL Category** tabs to define matching parameters for identifying traffic.<br>For example, select the **Application** tab, click **Add** and select web-browsing to apply the QoS Policy Rule to that application:<br> <br>(Optional) Define additional parameters. For example, on the **Source** tab, click **Add** to limit a specific user's web-browsing, in this case, user1:<br> 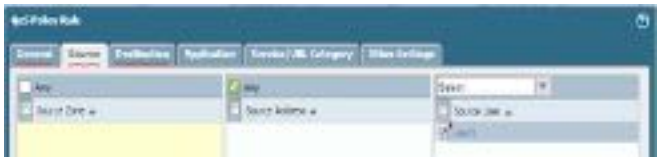 <br>4. On the **Other Settings** tab, select a QoS Class to assign to the QoS Policy Rule. For example, assign Class 2 to the user1's web-browsing traffic:<br> <br>5. Click **OK** to save the QoS Policy Rule. |


© Palo Alto Networks, Inc.

**Configure QoS (Continued)**

Step 5 Enable the QoS Profile on a physical interface.

You can configure settings to select cleartext and tunneled traffic for unique QoS treatment, in addition to the QoS configuration on the physical interface:

To configure specific settings for cleartext traffic using traffic's source interface and source subnet as criteria for QoS identification and treatment, perform Step 5 - 4.

To apply a QoS profile to a specific tunnel interface(s), perform Step 5 - 5

For more information, see QoS Cleartext and Tunneled Traffic.

Check if the platform you're using supports enabling QoS on a sub-interface by reviewing a summary of the Product Specifications.

It is a best practice to always define the **Egress Max** value for a QoS interface.

1. Select **Network > QoS** and click **Add** to open the QoS Interface dialog.
2. Enable QoS on the physical interface:
   a. On the **Physical Interface** tab, select the **Interface Name** of the interface to apply the QoS Profile to.

      In the example, Ethernet 1/1 is the egress interface for web-browsing traffic (see Step 2).

   b. Select **Turn on QoS feature on this interface**.
3. On the **Physical Interface** tab, select a QoS profile to apply by default to all **Clear Text** traffic.

   (Optional) Use the Tunnel Interface field to apply a QoS profile by default to all tunneled traffic.

For example, enable QoS on ethernet 1/1 and apply the QoS Profile named Limit Web Browsing as the default QoS Profile for clear text traffic.

4. (Optional) On the **Clear Text Traffic** tab, configure more granular QoS settings for cleartext traffic:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
   - Click **Add** to apply a QoS Profile to selected clear text traffic, further selecting the traffic for QoS treatment according to source interface and source subnet (creating a QoS node).
5. (Optional) On the **Tunneled Traffic** tab, configure more granular QoS settings for tunnel interfaces:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
   - Click **Add** to associate a selected tunnel interface with a QoS Profile.
6. Click **OK** to save the QoS Profile.
7. **Commit** the changes to enable the QoS Profile on the interface.

| Configure QoS (Continued) | |
|---|---|
| Step 6 Verify QoS configuration. | Select **Network > QoS** to view the **QoS Policies** page and click the **Statistics** link to view QoS bandwidth, active sessions of a selected QoS node or class, and active applications for the selected QoS node or class.<br>For example, see the statistics for ethernet 1/1 with QoS enabled:<br> 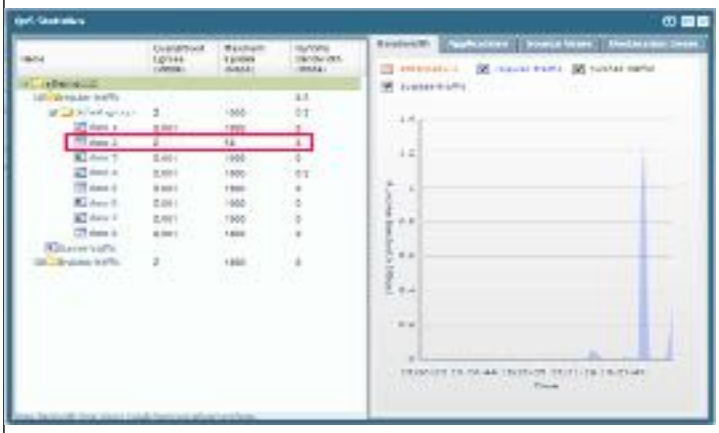 <br>Class 2 traffic limited to 2 Mbps of guaranteed bandwidth and a maximum bandwidth of 50 Mbps.<br>Continue to click the tabs to display further information regarding applications, source users, destination users, security rules and QoS rules.<br>Bandwidth limits shown on the **QoS Statistics** window include a hardware adjustment factor. |

 © Palo Alto Networks, Inc.

# Configure QoS for a Virtual System

QoS can be configured for a single or several Virtual Systems within a Palo Alto Networks firewall. Because a virtual system is an independent firewall, QoS must be configured independently for a single virtual system to apply a QoS configuration to only that virtual system.

Configuring QoS for a virtual system is similar to configuring QoS on a physical firewall, with the exception that configuring QoS for a virtual system requires specifying the traffic flow's source and destination zones and source and destination interfaces. Because a virtual system exists without set physical boundaries (such as a physical interface) that traffic flows through, specifying source and destination zones and interfaces of a traffic flow allows you to control and shape traffic for that virtual system specifically, as a traffic flow spans more than one virtual system in a virtual environment.

The example below shows two virtual systems configured within a firewall. VSYS 1 (purple) and VSYS 2 (red) each have QoS configured to prioritize or limit two distinct traffic flows, indicated by their corresponding purple (VSYS 1) and red (VSYS 2) lines. The QoS nodes indicate the points at which QoS traffic is identified and then shaped in each virtual system.

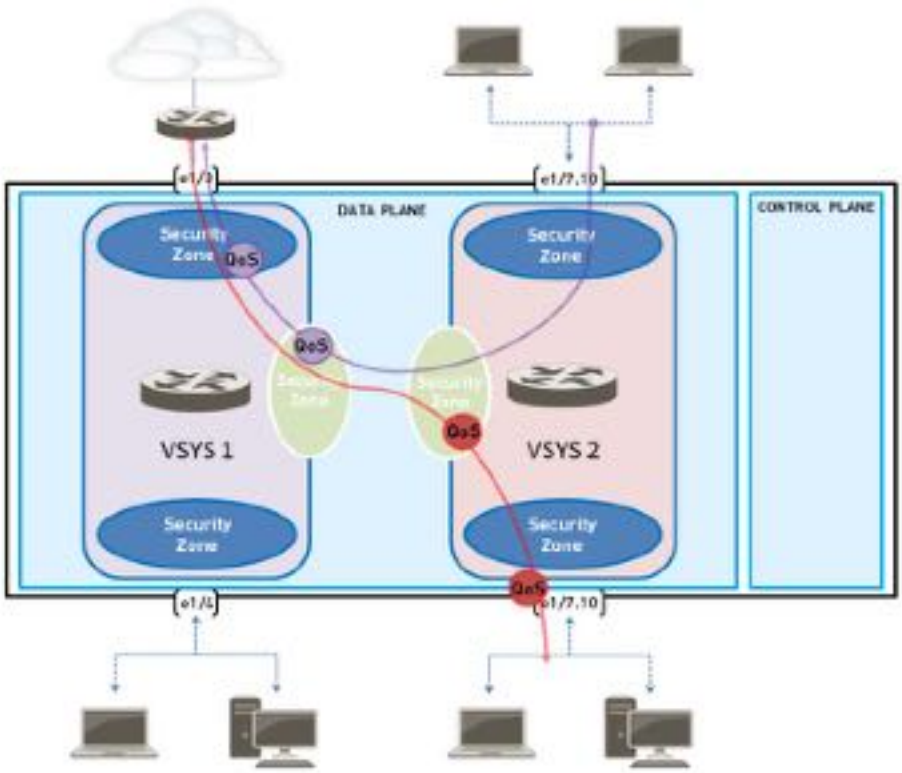

| Configure QoS in a Virtual System Environment | | |
|---|---|---|
| Step 1 | Confirm that the appropriate interfaces, virtual routers, and security zones are associated with each virtual system. | • To view configured interfaces, select **Network > Interface**.<br>• To view configured zones, select **Network > Zones**.<br>• To view information on defined virtual routers, select **Network > Virtual Routers**. |
| Step 2 | Identify traffic to apply QoS to. | Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.<br>To view information for a specific virtual system, select the virtual system from the **Virtual System** drop-down:<br> 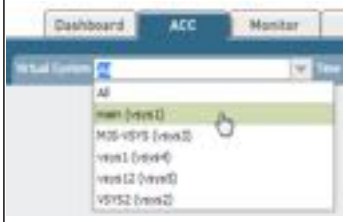 <br>Click any application name to display detailed application information. |


© Palo Alto Networks, Inc.

**Configure QoS in a Virtual System Environment**

Step 3 Identify the egress interface for applications that you identified as needing QoS treatment.

In a virtual system environment, QoS is applied to traffic on the traffic's egress point on the virtual system. Depending on a virtual system's configuration and the QoS policy, the egress point of QoS traffic could be associated with a physical interface or could be a configured zone.

This example shows how to limit web-browsing traffic on vsys 1.

Select **Monitor > Logs > Traffic** to view the device's traffic logs. Each entry has the option to display columns with information necessary to configure QoS in a virtual system environment:

- virtual system
- egress interface
- ingress interface
- source zone
- destination zone

To display a column if it is not displayed by default:

- Click any column header to add a column to the log:

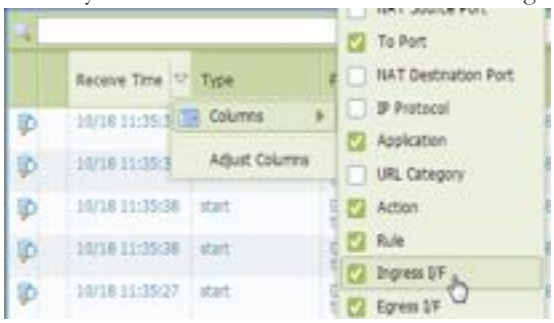


- Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface, as well as source and destination zones, in the **Source** and **Destination** sections:

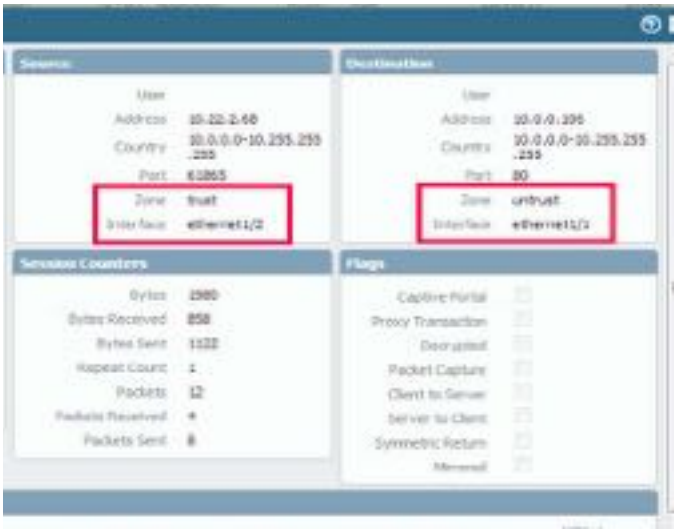

For example, for web-browsing traffic from VSYS 1, the ingress interface is ethernet 1/2, the egress interface is ethernet 1/1, the source zone is *trust* and the destination zone is *untrust*.

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 4 Create a QoS Profile.<br><br>You can edit any existing QoS Profile, including the default, by clicking the profile name. | 1. Select **Network > Network Profiles > QoS Profile** and click **Add** to open the QoS Profile dialog.<br>2. Enter a descriptive **Profile Name**.<br>3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS profile.<br>4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS profile.<br><br> Any traffic that exceeds the QoS profile's egress guaranteed limit is best effort but is not guaranteed.<br>5. In the Classes section of the **QoS Profile**, specify how to treat up to eight individual QoS classes:<br>a. Click **Add** to add a class to the QoS Profile.<br>b. Select the **Priority** for the class.<br>c. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.<br>d. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.<br>6. Click **OK** to save the QoS profile. |

 © Palo Alto Networks, Inc.

**Configure QoS in a Virtual System Environment**

Step 5 Create a QoS policy.

In a multi-VSYS environment, traffic can span more than one virtual system before the ingress point on the virtual system you are configuring QoS for. Specifying source and destination zones for QoS traffic ensures that the traffic is correctly identified as it flows through the specific virtual system (in this example, vsys 1) and QoS is applied to the traffic only for that designated virtual system (and not applied to traffic for other configured virtual systems).

1. Select **Policies > QoS** and click **Add** to open the QoS Policy Rule dialog.
2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**.
3. Specify the traffic to which the QoS policy rule will apply. Use the **Source**, **Destination**, **Application**, and **Service/URL Category** tabs to define matching parameters for identifying traffic.

   For example, select the **Application** tab, click **Add** and select web-browsing to apply the QoS Policy Rule to that application:

4. On the **Source** tab, click **Add** to select the source zone of vsys 1's web-browsing traffic.

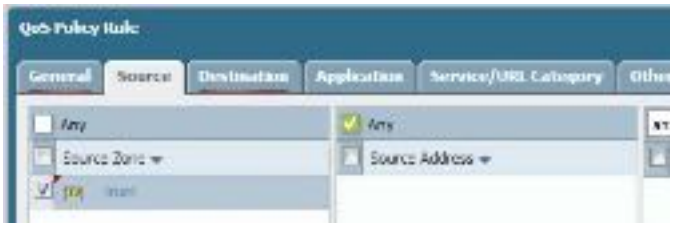


5. On the **Destination** tab, click **Add** to select the destination zone of vsys 1's web-browsing traffic.

6. On the **Other Settings** tab, and select a **QoS Class** to assign to the QoS policy rule. For example, assign Class 2 to web-browsing traffic on vsys 1:

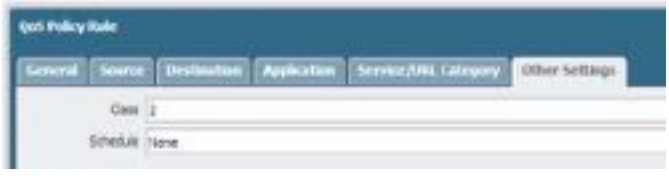


7. Click **OK** to save the QoS policy rule.

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 6 Enable the QoS Profile on a physical interface.<br><br>It is a best practice to always define the **Egress Max** value for a QoS interface. | 1. Select **Network > QoS** and click **Add** to open the QoS Interface dialog.<br>2. Enable QoS on the physical interface:<br>a. On the **Physical Interface** tab, select the **Interface Name** of the interface to apply the QoS Profile to.<br>In this example, ethernet 1/1 is the egress interface for web-browsing traffic on vsys 1 (see Step 2).<br> <br>b. Select **Turn on QoS feature on this interface**.<br>3. On the **Physical Interface** tab, select the default QoS profile to apply to all **Clear Text** traffic.<br>(Optional) Use the **Tunnel Interface** field to apply a default QoS profile to all tunneled traffic.<br>4. (Optional) On the **Clear Text Traffic** tab, configure additional QoS settings for clear text traffic:<br>• Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.<br>• Click **Add** to apply a QoS Profile to selected clear text traffic, further selecting the traffic for QoS treatment according to source interface and source subnet (creating a QoS node).<br>5. (Optional) On the **Tunneled Traffic** tab, configure additional QoS settings for tunnel interfaces:<br>• Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.<br>• Click **Add** to associate a selected tunnel interface with a QoS Profile.<br>6. Click **OK** to save changes.<br>7. **Commit** the changes. |


© Palo Alto Networks, Inc.

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 7 Verify QoS configuration. | • Select **Network > QoS** to view the QoS Policies page. The **QoS Policies** page verifies that QoS is enabled and includes a **Statistics** link. Click the Statistics link to view QoS bandwidth, active sessions of a selected QoS node or class, and active applications for the selected QoS node or class.<br>• In a multi-VSYS environment, sessions cannot span multiple systems. Multiple sessions are created for one traffic flow if the traffic passes through more than one virtual system. To browse sessions running on the firewall and view applied QoS Rules and QoS Classes, select **Monitor > Session Browser**. |

# QoS Use Case Examples

The following use cases demonstrate how to use QoS in common scenarios:

- Use Case: QoS for a Single User
- QoS for Voice and Video Applications


© Palo Alto Networks, Inc.

## Use Case: QoS for a Single User

A CEO finds that during periods of high network usage, she is unable to access enterprise applications to respond effectively to critical business communications. The IT admin wants to ensure that all traffic to and from the CEO receives preferential treatment over other employee traffic so that she is guaranteed not only access to, but high performance of, critical network resources.

### Apply QoS to a Single User

Step 1 The admin creates the QoS profile *CEO_traffic* to define how traffic originating from the CEO will be treated and shaped as it flows out of the company network:

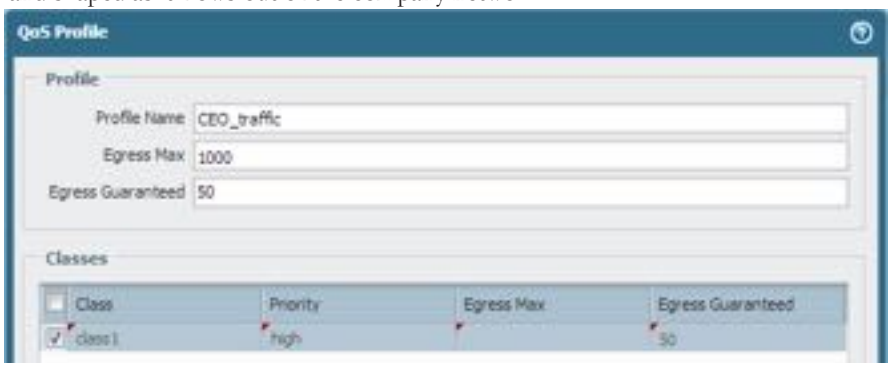


The admin assigns a guaranteed bandwidth (**Egress Guaranteed**) of 50 Mbps to ensure that the CEO will have that amount that bandwidth guaranteed to her at all times (more than she would need to use), regardless of network congestion.

The admin continues by designating Class 1 traffic as high priority and sets the profile's maximum bandwidth usage (**Egress Max**) to 1000 Mbps, the same maximum bandwidth for the interface that the admin will enable QoS on. The admin is choosing to not restrict the CEO's bandwidth usage in any way.

It is a best practice to populate the **Egress Max** field for a QoS profile, even if the max bandwidth of the profile matches the max bandwidth of the interface. The QoS profile's max bandwidth should never exceed the max bandwidth of the interface you are planning to enable QoS on.

**Apply QoS to a Single User (Continued)**

Step 2 The admin creates a QoS policy to identify the CEO's traffic (**Policies** > **QoS**) and assigns it the class that he defined in the QoS profile (see Step 1). Because User-ID is configured, the admin uses the **Source** tab in the QoS policy to singularly identify the CEO's traffic by her company network username. (If User-ID is not configured, the administrator could **Add** the CEO's IP address under **Source Address**. See User-ID.):

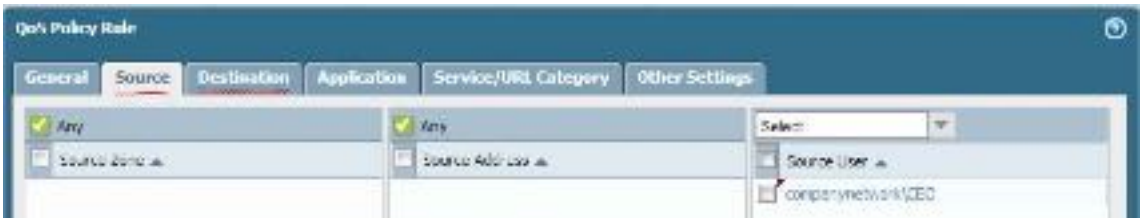


The admin associates the CEO's traffic with Class 1 (**Other Settings** tab) and then continues to populate the remaining required policy fields; the admin gives the policy a descriptive **Name** (**General** tab) and selects **Any** for the **Source Zone** (**Source** tab) and **Destination Zone** (**Destination** tab):

Step 3 Now that Class 1 is associated with the CEO's traffic, the admin enables QoS by checking **Turn on QoS feature on interface** and selecting the traffic flow's egress interface. The egress interface for the CEO's traffic flow is the external-facing interface, in this case, ethernet 1/2:

Because the admin wants to ensure that all traffic originating from the CEO is guaranteed by the QoS profile and associated QoS policy he created, he selects the *CEO_traffic* to apply to **Clear Text** traffic flowing from ethernet 1/2.

 © Palo Alto Networks, Inc.

**Apply QoS to a Single User (Continued)**

Step 4 After committing the QoS configuration, the admin navigates to the **Network > QoS** page to confirm that the QoS profile *CEO_traffic* is enabled on the external-facing interface, ethernet 1/2:

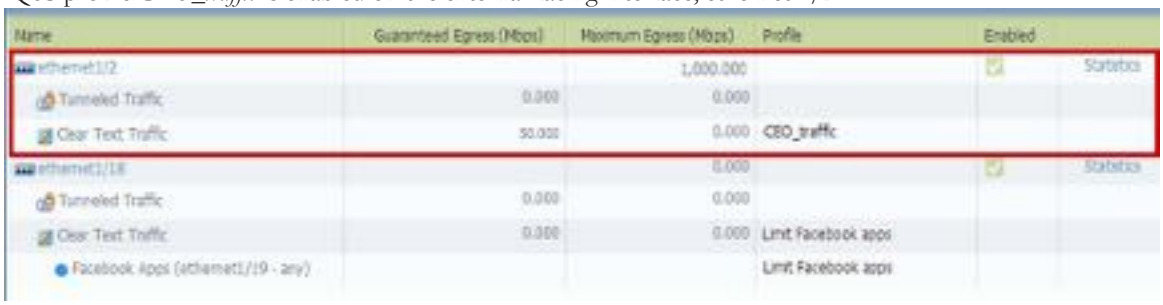


He clicks **Statistics** to view how traffic originating with the CEO (Class 1) is being shaped as it flows from ethernet 1/2:

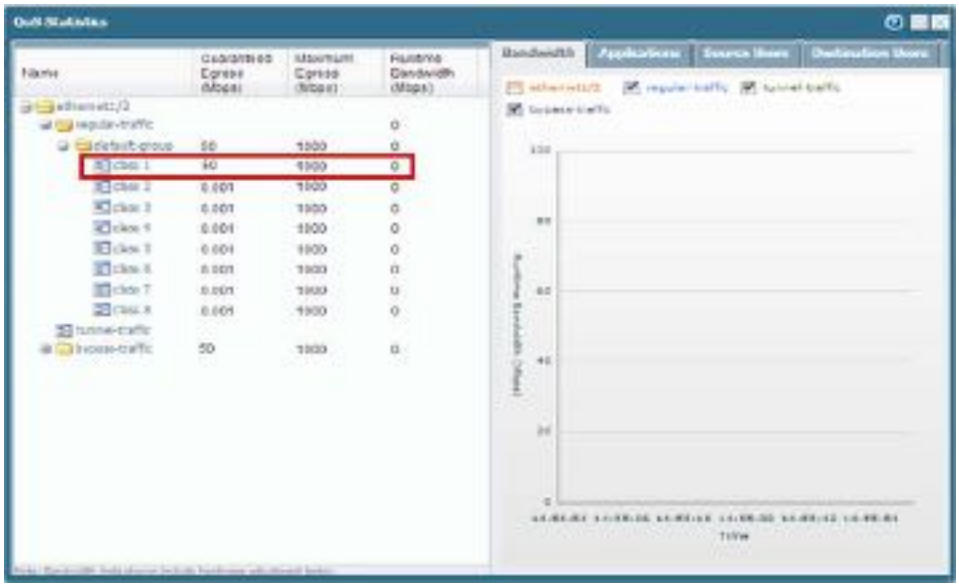

This case demonstrates how to apply QoS to traffic originating from a single source user. However, if you also wanted to guarantee or shape traffic to a destination user, you could configure a similar QoS setup. Instead of, or in addition to this work flow, create a QoS policy that specifies the user's IP address as the **Destination Address** on the **Policies > QoS** page (instead of specifying the user's source information, as shown in Step 2) and then enable QoS on the network's internal-facing interface on the **Network > QoS** page (instead of the external-facing interface, as shown in Step 3.)

## QoS for Voice and Video Applications

Voice and video traffic is particularly sensitive to measurements that the QoS feature shapes and controls, especially latency and jitter. For voice and video transmissions to be audible and clear, voice and video packets cannot be dropped, delayed, or delivered inconsistently. A best practice for voice and video applications, in addition to guaranteeing bandwidth, is to guarantee priority to voice and video traffic.

In this example, employees at a company branch office are experiencing difficulties and unreliability in using video conferencing and Voice-over-IP (VoIP) technologies to conduct business communications with other branch offices, with partners, and with customers. An IT admin intends to implement QoS in order to address these issues and ensure effective and reliable business communication for the branch employees. Because the admin wants to guarantee QoS to both incoming and outgoing network traffic, he will enable QoS on both the firewall's internal- and external-facing interfaces.

**Ensure Quality for Voice and Video Applications**

Step 1 The admin creates a QoS profile, defining Class 2 so that any traffic associated with Class 2 receives real-time priority and on an interface with a maximum bandwidth of 1000 Mbps, is guaranteed a bandwidth of 250 Mbps at all times, including peak periods of network usage.

Real-time priority is typically recommended for applications affected by latency, and is particularly useful in guaranteeing performance and quality of voice and video applications.

On the **Network > Network Profiles > Qos Profile** page, the admin clicks **Add**, enters the **Profile Name** *ensure voip-video traffic* and defines Class 2 traffic.

 © Palo Alto Networks, Inc.

**Ensure Quality for Voice and Video Applications (Continued)**

Step 2 The admin creates a QoS policy to identify voice and video traffic. Because the company does not have one standard voice and video application, the admin wants to ensure QoS is applied to a few applications that are widely and regularly used by employees to communicate with other offices, with partners, and with customers. On the **Policies > QoS > QoS Policy Rule > Applications** tab, the admin clicks **Add** and opens the **Application Filter** window. The admin continues by selecting criteria to filter the applications he wants to apply QoS to, choosing the Subcategory *voip-video*, and narrowing that down by specifying only voip-video applications that are both low-risk and widely-used.

The application filter is a dynamic tool that, when used to filter applications in the QoS policy, allows QoS to be applied to all applications that meet the criteria of *voip-video*, *low risk*, and *widely used* at any given time.

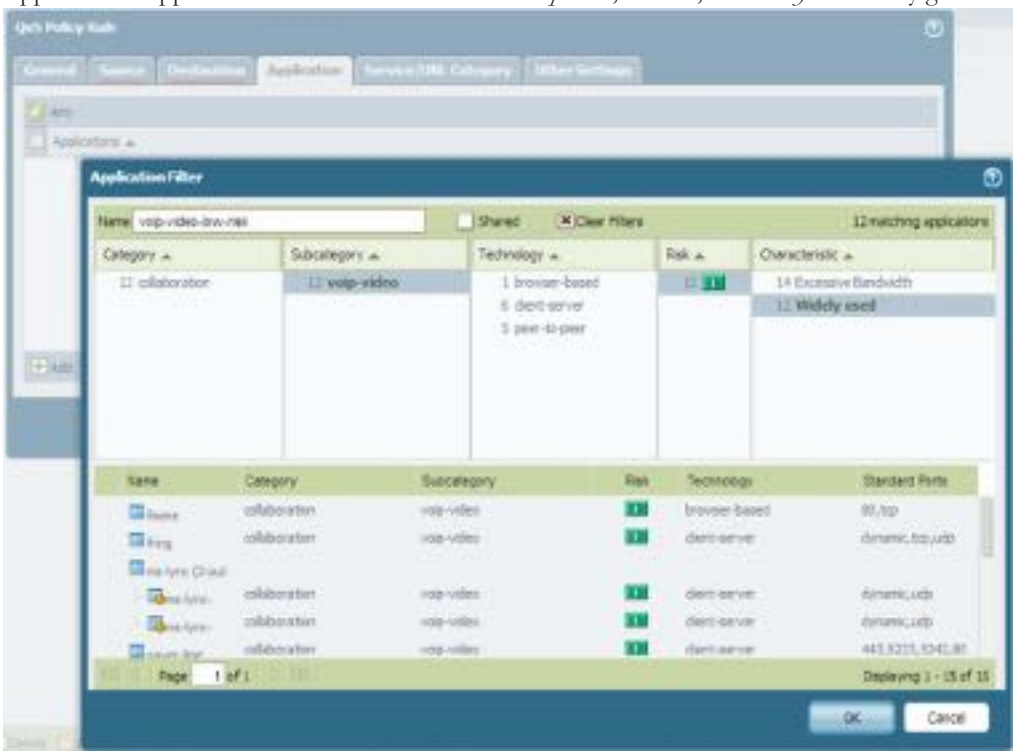


The admin names the **Application Filter** *voip-video-low-risk* and includes it in the QoS policy:

The admin names the QoS policy *Voice-Video* and associates the *voip-video-low-risk* application filter with Class 2 traffic (as he defined it in Step 1). He is going to use the *Voice-Video* QoS policy for both incoming and outgoing QoS traffic, so he sets **Source** and **Destination** information to **Any**:

**Ensure Quality for Voice and Video Applications (Continued)**

Step 3 Because the admin wants to ensure QoS for both incoming and outgoing voice and video communications, he enables QoS on the network's external-facing interface (to apply QoS to outgoing communications) and to the internal-facing interface (to apply QoS to incoming communications).

The admin begins by enabling the QoS profile he created in Step 1, *ensure voice-video traffic* (Class 1 in this profile is associated with policy created in Step 2, *Voice-Video*) on the external-facing interface, in this case, ethernet 1/2.

He then enables the same QoS profile *ensure voip-video traffic* on the internal-facing interface, in this case, ethernet 1/1.

Step 4 The admin confirms that QoS is enabled for both incoming and outgoing voice and video traffic:

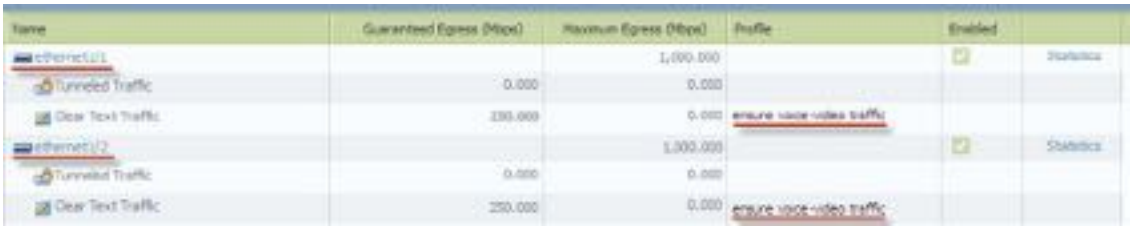


The admin has successfully enabled QoS on both the network's internal- and external-facing interfaces. Real-time priority is now ensured for voice and video application traffic as it flows both into and out of the network, ensuring that these communications, which are particularly sensitive to latency and jitter, can be used reliably and effectively to perform both internal and external business communications.

 © Palo Alto Networks, Inc.

# VPNs

Virtual private networks (VPNs) create tunnels that allow users/systems to connect securely over a public network, as if they were connecting over a local area network (LAN). To set up a VPN tunnel, you need a pair of devices that can authenticate each other and encrypt the flow of information between them. The devices can be a pair of Palo Alto Networks firewalls, or a Palo Alto Networks firewall along with a VPN-capable device from another vendor.


- VPN Deployments
- Site-to-Site VPN Overview
- Site-to-Site VPN Concepts
- Set Up Site-to-Site VPN
- Site-to-Site VPN Quick Configs


© Palo Alto Networks, Inc.

# VPN Deployments

The Palo Alto Networks firewall supports the following VPN deployments:

- **Site-to-Site VPN**— A simple VPN that connects a central site and a remote site, or a hub and spoke VPN that connects a central site with multiple remote sites. The firewall uses the IP Security (IPSec) set of protocols to set up a secure tunnel for the traffic between the two sites. See Site-to-Site VPN Overview.
- **Remote User-to-Site VPN**—A solution that uses the GlobalProtect agent to allow a remote user to establish a secure connection through the firewall. This solution uses SSL and IPSec to establish a secure connection between the user and the site. Refer to the GlobalProtect Administrator's Guide.
- **Large Scale VPN**— The Palo Alto Networks GlobalProtect Large Scale VPN (LSVPN) provides a simplified mechanism to roll out a scalable hub and spoke VPN with up to 1024 satellite offices. The solution requires Palo Alto Networks firewalls to be deployed at the hub and at every spoke. It uses certificates for device authentication, SSL for securing communication between all components, and IPSec to secure data. See Large Scale VPN (LSVPN).

**Figure: VPN Deployments**

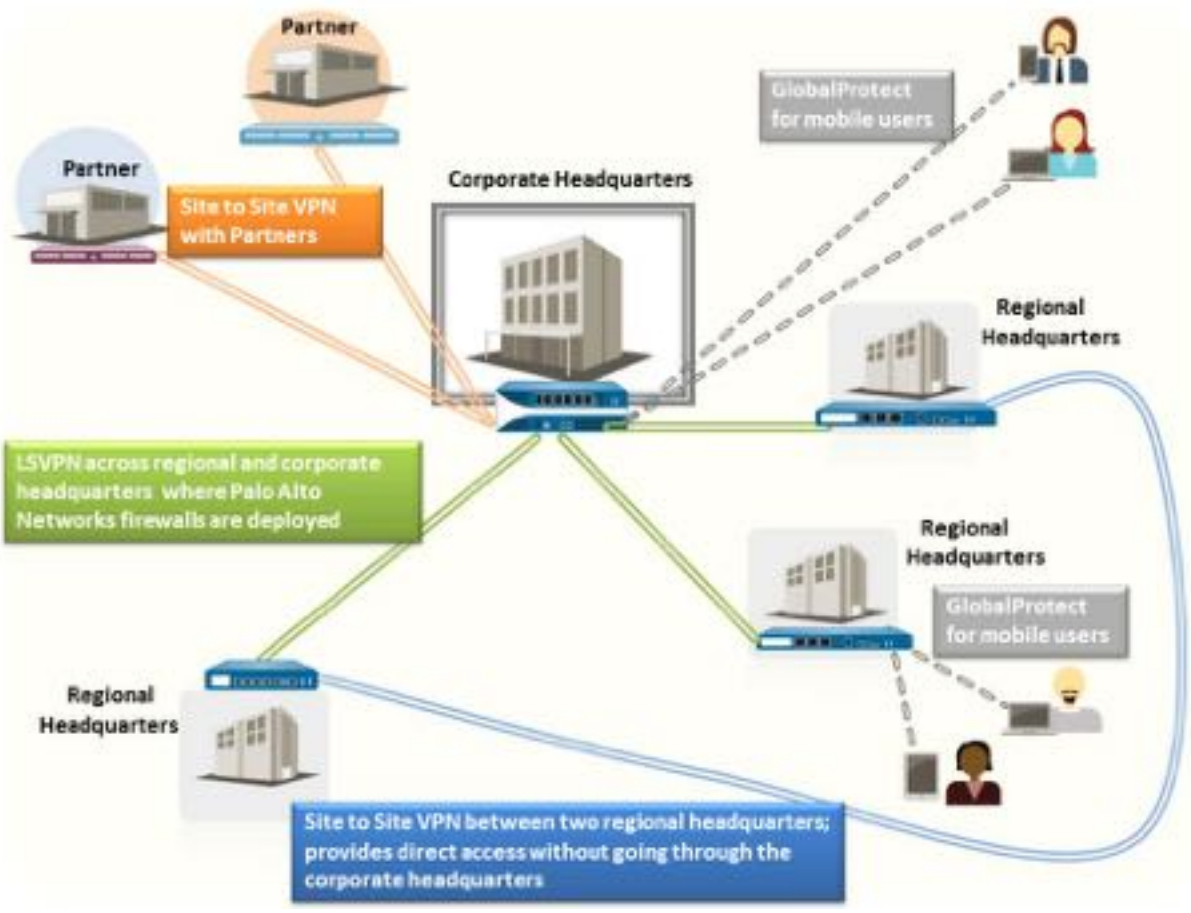


 © Palo Alto Networks, Inc.

# Site-to-Site VPN Overview

A VPN connection that allows you to connect two Local Area Networks (LANs) is called a site-to-site VPN. You can configure route-based VPNs to connect Palo Alto Networks firewalls located at two sites or to connect a Palo Alto Networks firewall with a third-party security device at another location. The firewall can also interoperate with third-party policy-based VPN devices; the Palo Alto Networks firewall supports route-based VPN.

The Palo Alto Networks firewall sets up a route-based VPN, where the firewall makes a routing decision based on the destination IP address. If traffic is routed to a specific destination through a VPN tunnel, then it is handled as VPN traffic.

The IP Security (IPSec) set of protocols is used to set up a secure tunnel for the VPN traffic, and the information in the TCP/IP packet is secured (and encrypted if the tunnel type is ESP). The IP packet (header and payload) is embedded in another IP payload, and a new header is applied and then sent through the IPSec tunnel. The source IP address in the new header is that of the local VPN peer and the destination IP address is that of the VPN peer on the far end of the tunnel. When the packet reaches the remote VPN peer (the firewall at the far end of the tunnel), the outer header is removed and the original packet is sent to its destination.

In order to set up the VPN tunnel, first the peers need to be authenticated. After successful authentication, the peers negotiate the encryption mechanism and algorithms to secure the communication. The Internet Key Exchange (IKE) process is used to authenticate the VPN peers, and IPSec Security Associations (SAs) are defined at each end of the tunnel to secure the VPN communication. IKE uses digital certificates or preshared keys, and the Diffie Hellman keys to set up the SAs for the IPSec tunnel.The SAs specify all of the parameters that are required for secure transmission— including the security parameter index (SPI), security protocol, cryptographic keys, and the destination IP address—encryption, data authentication, data integrity, and endpoint authentication.

The following figure shows a VPN tunnel between two sites. When a client that is secured by VPN Peer A needs content from a server located at the other site, VPN Peer A initiates a connection request to VPN Peer B. If the security policy permits the connection, VPN Peer A uses the IKE Crypto profile parameters (IKE phase 1) to establish a secure connection and authenticate VPN Peer B. Then, VPN Peer A establishes the VPN tunnel using the IPSec Crypto profile, which defines the IKE phase 2 parameters to allow the secure transfer of data between the two sites.

**Figure: Site-to-Site VPN**

# Site-to-Site VPN Concepts

A VPN connection provides secure access to information between two or more sites. In order to provide secure access to resources and reliable connectivity, a VPN connection needs the following components:

- IKE Gateway
- Tunnel Interface
- Tunnel Monitoring
- Internet Key Exchange (IKE) for VPN

## IKE Gateway

The Palo Alto Networks firewalls or a firewall and another security device that initiate and terminate VPN connections across the two networks are called the IKE Gateways. To set up the VPN tunnel and send traffic between the IKE Gateways, each peer must have an IP address—static or dynamic—or FQDN. The VPN peers use preshared keys or certificates to mutually authenticate each other.

The peers must also negotiate the mode—main or aggressive—for setting up the VPN tunnel and the SA lifetime in IKE Phase 1. Main mode protects the identity of the peers and is more secure because more packets are exchanged when setting up the tunnel. Main mode is the recommended mode for IKE negotiation if both peers support it. Aggressive mode uses fewer packets to set up the VPN tunnel and is hence faster but a less secure option for setting up the VPN tunnel.

See Set up an IKE Gateway for configuration details.

## Tunnel Interface

To set up a VPN tunnel, the Layer 3 interface at each end must have a logical *tunnel* interface for the firewall to connect to and establish a VPN tunnel. A tunnel interface is a logical (virtual) interface that is used to deliver traffic between two endpoints. Each tunnel interface can have a maximum of 10 IPSec tunnels; this means that up to 10 networks can be associated with the same tunnel interface on the firewall.

The tunnel interface must belong to a security zone to apply policy and it must be assigned to a virtual router in order to use the existing routing infrastructure. Ensure that the tunnel interface and the physical interface are assigned to the same virtual router so that the firewall can perform a route lookup and determine the appropriate tunnel to use.

Typically, the Layer 3 interface that the tunnel interface is attached to belongs to an external zone, for example the untrust zone. While the tunnel interface can be in the same security zone as the physical interface, for added security and better visibility, you can create a separate zone for the tunnel interface. If you create a separate zone for the tunnel interface, say a VPN zone, you will need to create security policies to enable traffic to flow between the VPN zone and the trust zone.

 © Palo Alto Networks, Inc.

To route traffic between the sites, a tunnel interface does not require an IP address. An IP address is only required if you want to enable tunnel monitoring or if you are using a dynamic routing protocol to route traffic across the tunnel. With dynamic routing, the tunnel IP address serves as the next hop IP address for routing traffic to the VPN tunnel.

If you are configuring the Palo Alto Networks firewall with a VPN peer that performs policy-based VPN, you must configure a local and remote Proxy ID when setting up the IPSec tunnel. Each peer compares the Proxy-IDs configured on it with what is actually received in the packet in order to allow a successful IKE phase 2 negotiation. If multiple tunnels are required, configure unique Proxy IDs for each tunnel interface; a tunnel interface can have a maximum of 250 ProxyIDs. Each Proxy ID counts towards the IPSec VPN tunnel capacity of the firewall, and the tunnel capacity varies by the firewall model.

See Set up an IPSec Tunnel for configuration details.

## Tunnel Monitoring

For a VPN tunnel, you can check connectivity to a destination IP address across the tunnel. The network monitoring profile on the firewall allows you to verify connectivity (using ICMP) to a destination IP address or a next hop at a specified polling interval, and to specify an action on failure to access the monitored IP address.

If the destination IP is unreachable, you either configure the firewall to wait for the tunnel to recover or configure automatic failover to another tunnel. In either case, the firewall generates a system log that alerts you to a tunnel failure and renegotiates the IPSec keys to accelerate recovery.

The default monitoring profile is configured to wait for the tunnel to recover; the polling interval is 3 seconds and the failure threshold is 5.

See Set up Tunnel Monitoring for configuration details.

## Internet Key Exchange (IKE) for VPN

The IKE process allows the VPN peers at both ends of the tunnel to encrypt and decrypt packets using mutually agreed-upon keys or certificate and method of encryption. The IKE process occurs in two phases: IKE Phase 1 and IKE Phase 2. Each of these phases use keys and encryption algorithms that are defined using cryptographic profiles— IKE crypto profile and IPSec crypto profile—and the result of the IKE negotiation is a Security Association (SA). An SA is a set of mutually agreed-upon keys and algorithms that are used by both VPN peers to allow the flow of data across the VPN tunnel. The following illustration depicts the key exchange process for setting up the VPN tunnel:

© Palo Alto Networks, Inc.

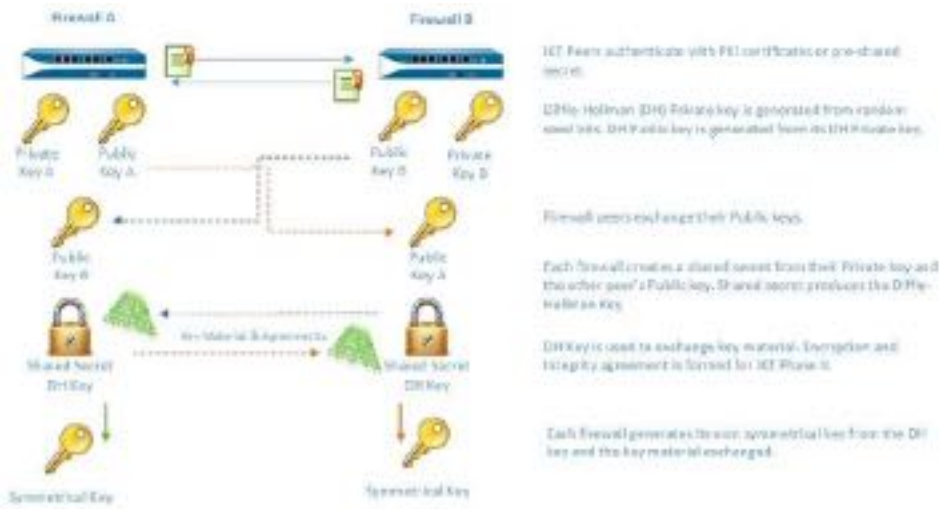

## IKE Phase 1

In this phase, the firewalls use the parameters defined in the IKE Gateway configuration and the IKE Crypto profile to authenticate each other and set up a secure control channel. IKE Phase supports the use of preshared keys or digital certificates (which use public key infrastructure, PKI) for mutual authentication of the VPN peers. Preshared keys are a simple solution for securing smaller networks because they do not require the support of a PKI infrastructure. Digital certificates can be more convenient for larger networks or implementations that require stronger authentication security.

When using certificates, make sure that the CA issuing the certificate is trusted by both gateway peers and that the maximum length of certificates in the certificate chain is 5 or less. With IKE fragmentation enabled, the firewall can reassemble IKE messages with up to 5 certificates in the certificate chain and successfully establish a VPN tunnel.

The IKE-crypto profile defines the following options that are used in the IKE SA negotiation:

- Diffie-Hellman (DH) Group for generating symmetrical keys for IKE. The Diffie Hellman algorithm uses the private key of one party and the public key of the other to create a shared secret, which is an encrypted key that is shared by both VPN tunnel peers. The DH groups supported on the firewall are: Group 1—768 bits; Group 2—1024 bits (the default); Group 5—1536 bits; Group 14—2048 bits.
- Authentication options—sha1; sha 256; sha 384; sha 512; md5
- Encryption algorithms—3des; aes128; aes192; aes256

## IKE Phase 2

After the tunnel is secured and authenticated, in Phase 2 the channel is further secured for the transfer of data between the networks. IKE Phase 2 uses the keys that were established in Phase 1 of the process and the IPSec Crypto profile, which defines the IPSec protocols and keys used for the SA in IKE Phase 2.

The IPSEC uses the following protocols to enable secure communication:

 © Palo Alto Networks, Inc.

- Encapsulating Security Payload (ESP)—Allows you to encrypt the entire IP packet, and authenticate the source and verify integrity of the data. While ESP requires that you encrypt and authenticate the packet, you can choose to only encrypt or only authenticate by setting the encryption option to Null; using encryption without authentication is discouraged.
- Authentication Header (AH)—Authenticates the source of the packet and verifies data integrity. AH does not encrypt the data payload and is unsuited for deployments where data privacy is important. AH is commonly used when the main concern is to verify the legitimacy of the peer, and data privacy is not required.

**Table: Algorithms Supported for IPSEC Authentication and Encryption**

| ESP | AH |
|---|---|
| Diffie Hellman Exchange options supported | |
| • Group 1—768 bits<br>• Group 2—1024 bits (the default)<br>• Group 5—1536 bits<br>• Group 14—2048 bits.<br>• no-pfs—By default, Perfect Forward Secrecy (pfs) is enabled. With PFS is enabled, a new DH key is generated in IKE phase 2 using one of the groups listed above; this key is independent of the keys exchanged in IKE phase1, and therefore allows for more secure transfer of data.<br>No-pfs implies that the DH key created at phase 1 is not renewed and a single key is used for the IPSEC SA negotiations. Both VPN peers must be enabled or disabled for Perfect Forward Secrecy. | |
| Encryption algorithms supported | |
| • 3des | |
| • aes128 | |
| • aes192 | |
| • aes256 | |
| • aes128ccm16 | |
| • null | |
| Authentication algorithms supported | |
| • md5 | • md5 |
| • sha 1 | • sha 1 |
| • sha 256 | • sha 256 |
| • sha 384 | • sha 384 |
| • sha512 | • sha 512 |
| • none | |

## Methods of Securing IPSec VPN Tunnels (IKE Phase 2)

IPSec VPN tunnels can be secured using manual keys or auto keys. In addition, IPSec configuration options include Diffie-Hellman Group for key agreement, and/or an encryption algorithm and a hash for message authentication.

- **Manual Key**—Manual key is typically used if the Palo Alto Networks firewall is establishing a VPN tunnel with a legacy device, or if you want to reduce the overhead of generating session keys. If using manual keys, the same key must be configured on both peers.

  Manual keys are not recommended for establishing a VPN tunnel because the session keys can be compromised when relaying the key information between the peers; if the keys are compromised, the data transfer is no longer secure.

- **Auto Key**— Auto Key allows you to automatically generate keys for setting up and maintaining the IPSec tunnel based on the algorithms defined in the IPSec Crypto profile.

 © Palo Alto Networks, Inc.

# Set Up Site-to-Site VPN

To set up site-to-site VPN:

- Make sure that your Ethernet interfaces, virtual routers, and zones are configured properly. For more information, see Set Up Interfaces and Zones.
- Create your tunnel interfaces. Ideally, put the tunnel interfaces in a separate zone, so that tunneled traffic can use different policies.
- Set up static routes or assign routing protocols to redirect traffic to the VPN tunnels. To support dynamic routing (OSPF, BGP, RIP are supported), you must assign an IP address to the tunnel interface.
- Define IKE gateways for establishing communication between the peers across each end of the VPN tunnel; also define the cryptographic profile that specifies the protocols and algorithms for identification, authentication, and encryption to be used for setting up VPN tunnels in IKEv1 Phase 1. See Set up an IKE Gateway and Define IKE Crypto Profiles.
- Configure the parameters that are needed to establish the IPSec connection for transfer of data across the VPN tunnel; See Set up an IPSec Tunnel. For IKEv1 Phase-2, see Define IPSec Crypto Profiles.
- (Optional) Specify how the firewall will monitor the IPSec tunnels. See Set up Tunnel Monitoring.
- Define security policies to filter and inspect the traffic.

If there is a deny rule at the end of the security rulebase, intra-zone traffic is blocked unless otherwise allowed. Rules to allow IKE and IPSec applications must be explicitly included above the deny rule.

When these tasks are complete, the tunnel is ready for use. Traffic destined for the zones/addresses defined in policy is automatically routed properly based on the destination route in the routing table, and handled as VPN traffic. For a few examples on site-to-site VPN, see Site-to-Site VPN Quick Configs.

## Set up an IKE Gateway

To set up a VPN tunnel, the VPN peers or gateways must authenticate each other using preshared keys or digital certificates and establish a secure channel in which to negotiate the IPSec security association (SA) that will be used to secure traffic between the hosts on each side.

| **Set up an IKE Gateway** | |
|---|---|
| Step 1 Define the new IKE Gateway. | 1. Select **Network > Network Profiles > IKE Gateway** and enter a **Name f**or the new gateway configuration.<br>2. Select the outgoing **Interface** on the firewall.<br>3. From the **Local IP address** drop-down list, select the IP address that will be used to as the endpoint for the VPN connection. This is the external-facing interface with a publicly routable IP address on the firewall. |

| Set up an IKE Gateway | | |
|---|---|---|
| Step 2 | Define the settings for the peer at the far end of the tunnel. | 1. Select whether the peer uses a **Static** or **Dynamic** IP address in **Peer IP Type**.<br>2. If the **Peer IP Address** is static, enter the IP address of the peer. |
| Step 3 | Select the peer authentication method.<br>This is required for static and dynamic peers. | • For configuring a **Pre-Shared Key**, see Step 4.<br>• For configuring digital **Certificates**, see Step 5. |
| Step 4 | Configure a pre-shared key. | 1. Enter a security key to use for authentication across the tunnel. This key must be the same on both peers.<br>Generate a key that is hard to crack with dictionary attacks; use a pre-shared key generator, if necessary.<br>2. Continue with Step 6. |
| Step 5 | Configure certificate-based authentication.<br> The pre-requisites for certificate-based authentication are as follows:<br>– **Obtain a signed certificate:** See Generate a Certificate on the Firewall or Obtain a Certificate from an External CA.<br>– **Configure the certificate profile**: The certificate profile provides the settings that the IKE gateway uses for negotiating and validating certificate authentication with its peer. See Configure a Certificate Profile. | 1. Select **Certificate** for the **Authentication** method and select the signed certificate from the **Local Certificate** drop-down.<br>If your device is enabled for multi-virtual systems, if the certificate belongs to a virtual system, it must be in the same virtual system as the interface used for the IKE gateway.<br>2. From the **Local Identification** drop-down list, choose one of the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), Distinguished Name (subject).<br>3. From the **Peer Identification** drop-down, choose one of the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), Distinguished Name (subject).<br>4. Select a **Certificate Profile** to use.<br>5. Continue with Step 6. |

 © Palo Alto Networks, Inc.

| Set up an IKE Gateway | |
|---|---|
| Step 6 Configure the additional parameters for IKE phase 1 negotiations—Exchange mode, Crypto profile, IKE fragmentation, Dead Peer Detection. | 1. Select **Network > Network Profiles > IKE Gateways** and select the **Advanced Phase 1 Options** tab.<br>2. Choose **auto**, **aggressive**, or **main** for the **Exchange Mode**.<br>When a device is set to use the **auto** exchange mode, it can accept both main mode and aggressive mode negotiation requests; however, whenever possible, it initiates negotiation and allows exchanges in main mode.<br> If the exchange mode is not set to **auto**, you must configure both VPN peers with the same exchange mode to allow each peer to accept negotiation requests.<br>3. Select an existing profile or keep the default profile from **IKE Crypto Profile** drop-down. For details on defining an IKE Crypto profile, see Define IKE Crypto Profiles.<br>4. Select **Passive Mode** if you want the firewall to only respond to IKE connections and never initiate them.<br>5. Select **NAT Traversal Select** to have UDP encapsulation used on IKE and UDP protocols, enabling them to pass through intermediate NAT devices.<br>6. (Only if using certificate-based authentication and the exchange mode is not set as **aggressive** mode) Select **Enable Fragmentation** to enable the firewall to operate with IKE Fragmentation.<br>7. Select the **Dead Peer Detection** check box and enter an **Interval** (2 - 100 seconds); For **Retry,** define the time to delay (2 - 100 seconds) before attempting to re-check availability.<br>Dead peer detection identifies inactive or unavailable IKE peers by sending an IKE phase 1 notification payload to the peer and waiting for an acknowledgment. |
| Step 7 Save the changes. | Click **OK** and **Commit**. |

## Define Cryptographic Profiles

A cryptographic profile specifies the ciphers used for authentication and/or encryption between two IKE peers, and the lifetime of the key. The time period between each renegotiation is known as the lifetime; when the specified time expires, the firewall renegotiates a new set of keys.

For securing communication across the VPN tunnel, the firewall requires IKE and IPSec cryptographic profiles for completing IKE phase 1 and phase 2 negotiations, respectively. The firewall includes a *default* IKE crypto profile and a *default* IPSec crypto profile that is ready for use.

- Define IKE Crypto Profiles
- Define IPSec Crypto Profiles

## Define IKE Crypto Profiles

The IKE crypto profile is used to set up the encryption and authentication algorithms used for the key exchange process in IKE Phase 1, and lifetime of the keys, which specifies how long the keys are valid. To invoke the profile, you must attach it to the IKE Gateway configuration.

All IKE gateways configured on the same interface or local IP address must use the same crypto profile.

| Define an IKE Crypto Profile | | |
|---|---|---|
| Step 1 | Create a new IKE profile. | 1. Select **Network > Network Profiles > IKE Crypto** and select **Add**.<br>2. Enter a **Name f**or the new profile. |
| Step 2 | Select the DH Group to use for setting up the key exchange. | Click **Add** and select the key strength that you want to use for the DH Group.<br>Select more than one DH group if you are not certain of what is supported by the VPN peer. Prioritize the list of ciphers by strength so that the strongest cipher is used for setting up the tunnel. |
| Step 3 | Select the authentication and encryption algorithm. | Click **Add** and select the **Authentication** and **Encryption** algorithms that you want to use for communication between the IKE peers.<br>Selecting multiple algorithms allows the peers to use the strongest cipher/algorithm that is supported on both IKE peers. |
| Step 4 | Specify the duration for which the key is valid. | Select the **Lifetime** for which the key is valid. The time period between each renegotiation is known as the lifetime; when the specified time expires, the firewall will renegotiate a new set of keys. |
| Step 5 | Save your IKE Crypto profile. | Click **OK** and click **Commit**. |
| Step 6 | Attach the IKE Crypto profile to the IKE Gateway configuration. | See Step 6 in Set up an IKE Gateway. |

## Define IPSec Crypto Profiles

The IPSec crypto profile is invoked in IKE Phase 2. It specifies how the data is secured within the tunnel when Auto Key IKE is used to automatically generate keys for the IKE SAs.

 © Palo Alto Networks, Inc.

| Define the IPSec Crypto Profile | | |
|---|---|---|
| Step 1 | Create a new IPSec profile. | 1. Select **Network > Network Profiles > IPSec Crypto** and select **Add**.<br>2. Enter a **Name f**or the new profile.<br>3. Select the **IPSec Protocol**—ESP or AH—that you want to apply to secure the data as it traverses across the tunnel.<br>4. Click **Add** and select the **Authentication** and **Encryption** algorithms for ESP, and **Authentication** algorithms for AH, so that the IKE peers can negotiate the keys for the secure transfer of data across the tunnel.<br>Selecting multiple algorithms allows the peers to use the strongest cipher/algorithm that is supported on both IKE peers. |
| Step 2 | Select the DH Group to use for the IPSec SA negotiations in IKE phase 2. | 1. Select the key strength that you want to use from the **DH Group** drop-down.<br>Select more than one DH group if you are not certain of what key strength is supported by the peer at the other end. The strongest cipher will be used for setting up the tunnel.<br>2. Select **no-pfs**, if you do not want to renew the key that was created at phase 1, the current key is reused for the IPSEC SA negotiations. |
| Step 3 | Specify the duration of the key—time and volume of traffic. | Using a combination of time and traffic volume allows you to ensure safety of data.<br>Select the **Lifetime** or time period for which the key is valid. When the specified time expires, the firewall will renegotiate a new set of keys.<br>Select the **Lifesize** or volume of data after which the keys must be renegotiated. |
| Step 4 | Save your IPSec profile. | Click **OK** and click **Commit**. |
| Step 5 | Attach the IPSec Profile to an IPSec tunnel configuration. | See Step 4 in |

## Set up an IPSec Tunnel

The IPSec tunnel configuration allows you to authenticate and/or encrypt the data (IP packet) as it traverses across the tunnel.

If you are setting up the Palo Alto Networks firewall to work with a peer that supports policy-based VPN, you must define Proxy IDs. Devices that support policy-based VPN use specific security rules/policies or access-lists (source addresses, destination addresses and ports) for permitting interesting traffic through an IPSec tunnel. These rules are referenced during quick mode/IKE phase 2 negotiation, and are exchanged as Proxy-IDs in the first or the second message of the process. So, if you are configuring the Palo Alto Networks firewall to work with a policy-based VPN peer, for a successful phase 2 negotiation you must define the

Proxy-ID so that the setting on both peers is identical. If the Proxy-ID is not configured, because the Palo Alto Networks firewall supports route-based VPN, the default values used as Proxy-ID are source ip: 0.0.0.0/0, destination ip: 0.0.0.0/0 and application: any; and when these values are exchanged with the peer, it results in a failure to set up the VPN connection.

**Set up an IPSec Tunnel**

Step 1 Select **Network > IPSec Tunnels> General** and enter a **Name** for the new tunnel.

Step 2 Select the **Tunnel interface** that will be used to set up the IPSec tunnel.

– To create a new tunnel interface:

1. Select **Network > Interfaces > Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:

   • To use your trust zone as the termination point for the tunnel, select the zone from the drop-down. Associating the tunnel interface with the same zone (and virtual router) as the external-facing interface on which the packets enter the firewall, mitigates the need to create inter-zone routing.

   • (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example vpn-corp), and click **OK**.
4. In the **Virtual Router** drop-down, select **default**.
5. (Optional) If you want to assign an IPv4 address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.31.32.1/32.
6. If you want to assign an IPv6 address to the tunnel interface, see Step 3.
7. To save the interface configuration, click **OK**.

Step 3 (Optional) Enable IPv6 on the tunnel interface.

1. Select the IPv6 tab on **Network > Interfaces > Tunnel > IPv6**.
2. Select the check box to **Enable IPv6 on the interface**.

   This option allows you to route IPv6 traffic over an IPv4 IPSec tunnel and will provide confidentiality between IPv6 networks. The IPv6 traffic is encapsulated by IPv4 and then ESP. To route IPv6 traffic to the tunnel, you can use a static route to the tunnel, or use OSPFv3, or use a Policy-Based Forwarding (PBF) rule to direct traffic to the tunnel.
3. Enter the 64-bit extended unique **Interface ID** in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. By default, the firewall will use the EUI-64 generated from the physical interface's MAC address.
4. To enter an IPv6 **Address**, click **Add** and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. If Prefix is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.
   a. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address.
   b. Select **Anycast** to include routing through the nearest node.

 © Palo Alto Networks, Inc.

| Set up an IPSec Tunnel | |
|---|---|
| Step 4 Select the type of key that will be used to secure the IPSec tunnel. | Continue to one of the following steps, depending on why type of key exchange you are using:<br>• Set up Auto Key exchange.<br>• Set up Auto Key exchange. |
| • Set up Auto Key exchange. | 1. Select the IKE Gateway. To set up an IKE gateway, see Set up an IKE Gateway.<br>2. (Optional) Select the default IPSec Crypto Profile. To create a new IPSec Profile, see Define IPSec Crypto Profiles. |
| • Set up a Manual Key exchange. | 1. Set up the parameters for the local firewall:<br>a. Specify the **SPI** for the local firewall. SPI is a 32-bit hexadecimal index that is added to the header for IPSec tunneling to assist in differentiating between IPSec traffic flows; it is used to create the SA required for establishing a VPN tunnel.<br>b. Select the **Interface** that will be the tunnel endpoint, and optionally select the IP address for the local interface that is the endpoint of the tunnel.<br>c. Select the protocol to be used—**AH** or **ESP**.<br>d. For AH, select the **Authentication** method from the drop-down and enter a **Key** and then **Confirm Key**.<br>e. For ESP, select the **Authentication** method from the drop-down and enter a **Key** and then **Confirm Key**. Then, select the **Encryption** method and enter a **Key** and then **Confirm Key**, if needed.<br>2. Set up the parameters that pertain to the remote VPN peer.<br>a. Specify the **SPI** for the remote peer.<br>b. Enter the **Remote Address**, the IP address of the remote peer. |
| Step 5 Protect against a replay attack.<br>A replay attack occurs when a packet is maliciously intercepted and retransmitted by the interceptor. | Select the **Show Advanced Options** check box, select **Enable Replay Protection** to detect and neutralize against replay attacks. |
| Step 6 Preserve the Type of Service header for the priority or treatment of IP packets. | In the **Show Advanced Options** section, select **Copy TOS Header**. This copies the Type of Service (TOS) header from the inner IP header to the outer IP header of the encapsulated packets in order to preserve the original TOS information. |
| Step 7 Enable Tunnel Monitoring.<br> You need to assign an IP address to the tunnel interface for monitoring. | To alert the device administrator to tunnel failures and to provide automatic failover to another tunnel interface:<br>1. Specify a **Destination IP** address on the other side of the tunnel to determine if the tunnel is working properly.<br>2. Select a **Profile** to determine the action on tunnel failure. To create a new profile, see Define a Tunnel Monitoring Profile. |

| Set up an IPSec Tunnel | | |
|---|---|---|
| Step 8 | (Required only if the VPN peer uses policy-based VPN). Create a Proxy ID to identify the VPN peers. | 1. Select **Network > IPSec Tunnels > ProxyID**.<br>2. Click **Add** and enter the IP address for the VPN gateway peers.<br> 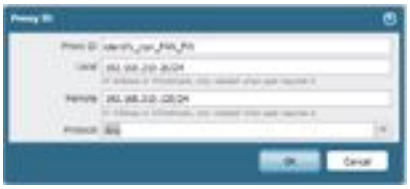 |
| Step 9 | Save your changes | Click **OK** and **Commit**. |

## Set up Tunnel Monitoring

To provide uninterrupted VPN service, you can use the Dead Peer Detection capability along with the tunnel monitoring capability on the firewall. You can also monitor the status of the tunnel. These monitoring tasks are described in the following sections:

- Define a Tunnel Monitoring Profile
- View the Status of the Tunnels

### Define a Tunnel Monitoring Profile

A tunnel monitoring profile allows you to verify connectivity between the VPN peers; you can configure the tunnel interface to ping a destination IP address at a specified interval and specify the action if the communication across the tunnel is broken.

| Define a Tunnel Monitoring Profile | |
|---|---|
| Step 1 | Select **Network > Network Profiles > Monitor**. A default tunnel monitoring profile is available for use. |
| Step 2 | Click **Add**, and enter a **Name** for the profile. |
| Step 3 | Select the **Action** if the destination IP address is unreachable.<br>• **Wait Recover**—the firewall waits for the tunnel to recover. It continues to use the tunnel interface in routing decisions as if the tunnel were still active.<br>• **Fail Over**—forces traffic to a back-up path if one is available. The firewall disables the tunnel interface, and thereby disables any routes in the routing table that use the interface.<br>In either case, the firewall attempts to accelerate the recovery by negotiating new IPSec keys. |
| Step 4 | Specify the **Interval** and **Threshold** to trigger the specified action.<br>The threshold specifies the number of heartbeats to wait before taking the specified action. The range is 2-100 and the default is 5.<br>The Interval measures the time between heartbeats. The range is 2-10 and the default is 3 seconds. |
| Step 5 | Attach the monitoring profile to the IPsec Tunnel configuration. See Enable Tunnel Monitoring. |

 © Palo Alto Networks, Inc.

## View the Status of the Tunnels

The status of the tunnel informs you about whether or not valid IKE phase-1 and phase-2 SAs have been established, and whether the tunnel interface is up and available for passing traffic.

Because the tunnel interface is a logical interface, it cannot indicate a physical link status. Therefore, you must enable tunnel monitoring so that the tunnel interface can verify connectivity to an IP address and determine if the path is still usable. If the IP address is unreachable, the firewall will either wait for the tunnel to recover or failover. When a failover occurs, the existing tunnel is torn down and routing changes are triggered to set up a new tunnel and redirect traffic.

**View Tunnel Status**

1. Select **Network > IPSec Tunnels**.
2. View the **Tunnel Status**.
   - Green indicates a valid IPSec SA tunnel.
   - Red indicates that IPSec SA is not available or has expired.
3. View the **IKE Gateway Status**.
   - Green indicates a valid IKE phase-1 SA.
   - Red indicates that IKE phase-1 SA is not available or has expired.
4. View the **Tunnel Interface Status**.
   - Green indicates that the tunnel interface is up.
   - Red indicates that the tunnel interface is down, because tunnel monitoring is enabled and the status is down.

To troubleshoot a VPN tunnel that is not yet up, see Interpret VPN Error Messages.

## Test VPN Connectivity

**Test Connectivity**

- Initiate IKE phase 1 by either pinging a host across the tunnel or using the following CLI command:

  **test vpn ike-sa gateway** *gateway_name*
- Then enter the following command to test if IKE phase 1 is set up:

  **show vpn ike-sa gateway** *gateway_name*

  In the output, check if the Security Association displays. If it does not, review the system log messages to interpret the reason for failure.
- Initiate IKE phase 2 by either pinging a host from across the tunnel or using the following CLI command:

  **test vpn ipsec-sa tunnel** *tunnel_name*
- Then enter the following command to test if IKE phase 1 is set up:

  **show vpn ipsec-sa tunnel** *tunnel_name*

  In the output, check if the Security Association displays. If it does not, review the system log messages to interpret the reason for failure.
- To view the VPN traffic flow information, use the following command:

  **show vpn-flow**

```
admin@PA-500> show vpn flow
total tunnels configured:             1
filter - type IPSec, state any

total IPSec tunnel configured:        1
total IPSec tunnel shown:             1

name               id      state    local-ip       peer-ip       tunnel-i/f
------------------------------------------------------------------------------
vpn-to-siteB       5       active    100.1.1.1     200.1.1.1      tunnel.41
```

## Interpret VPN Error Messages

The following table lists some of the common VPN error messages that are logged in the system log.

 © Palo Alto Networks, Inc.

**Table: Syslog Error Messages for VPN Issues**

| If error is this: | Try this: |
|---|---|
| `IKE phase-1 negotiation is failed as initiator, main mode. Failed SA: x.x.x.x[500]-y.y.y.y[500] cookie:84222f276c2fa2e9:0000000000000000 due to timeout.`<br>or<br>`IKE phase 1 negotiation is failed. Couldn't find configuration for IKE phase-1 request for peer IP x.x.x.x[1929]` | • Verify that the public IP address for each VPN peer is accurate in the IKE Gateway configuration.<br>• Verify that the IP addresses can be pinged and that routing issues are not causing the connection failure. |
| `Received unencrypted notify payload (no proposal chosen) from IP x.x.x.x[500] to y.y.y.y[500], ignored...`<br>or<br>`IKE phase-1 negotiation is failed. Unable to process peer's SA payload.` | Check the IKE Crypto profile configuration to verify that the proposals on both sides have a common encryption, authentication, and DH Group proposal. |
| `pfs group mismatched:my: 2peer: 0`<br>or<br>`IKE phase-2 negotiation failed when processing SA payload. No suitable proposal found in peer's SA payload.` | Check the IPSec Crypto profile configuration to verify that:<br>• pfs is either enabled or disabled on both VPN peers<br>• the DH Groups proposed by each peer has at least one DH Group in common |
| `IKE phase-2 negotiation failed when processing Proxy ID. Received local id x.x.x.x/x type IPv4 address protocol 0 port 0, received remote id y.y.y.y/y type IPv4 address protocol 0 port 0.` | The VPN peer on one end is using policy-based VPN. You must configure a Proxy ID on the Palo Alto Networks firewall. See Step 8. |

# Site-to-Site VPN Quick Configs

The following sections provide instructions for configuring some common VPN deployments:

- Site-to-Site VPN with Static Routing
- Site-to-Site VPN with OSPF
- Site-to-Site VPN with Static and Dynamic Routing

## Site-to-Site VPN with Static Routing

The following example shows a VPN connection between two sites that use static routes. Without dynamic routing, the tunnel interfaces on VPN Peer A and VPN Peer B do not require an IP address because the firewall automatically uses the tunnel interface as the next hop for routing traffic across the sites. However, to enable tunnel monitoring, a static IP address has been assigned to each tunnel interface.

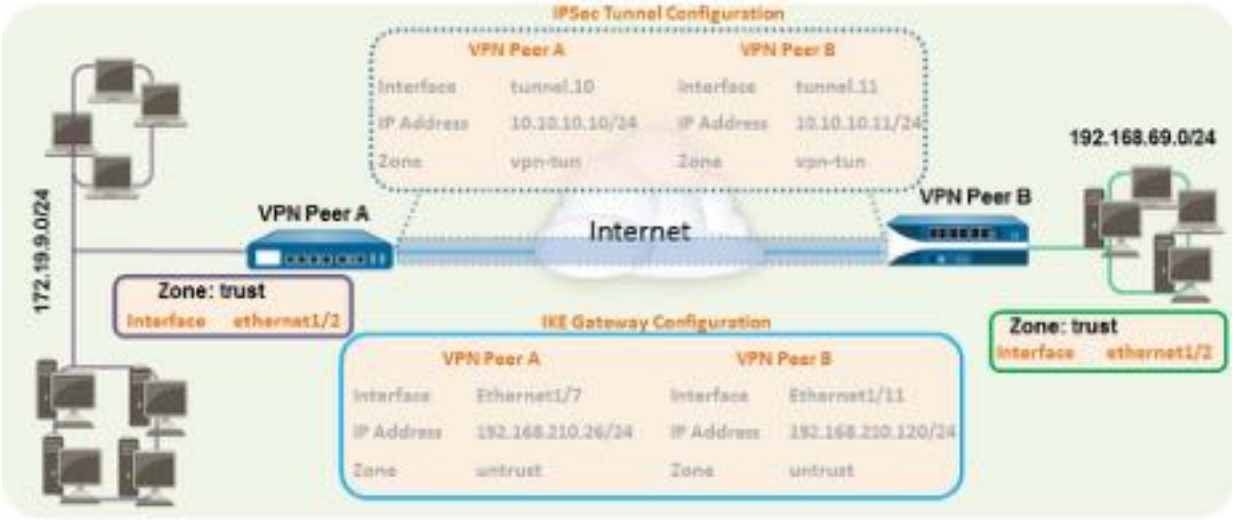


 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static Routing | |
|---|---|
| Step 1 Configure a Layer 3 interface.<br><br>This interface is used for the IKE phase-1 tunnel. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—192.168.210.26/24<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—192.168.210.120/24 |

| Quick Config: Site-to-Site VPN with Static Routing | |
|---|---|
| Step 2 Create a tunnel interface and attach it to a virtual router and security zone.<br> 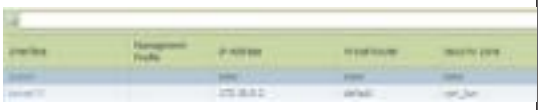 | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, such as **.1**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. (Optional) Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface.<br>With static routes, the tunnel interface does not require an IP address. For traffic that is destined to a specified subnet/IP address, the tunnel interface will automatically become the next hop. Consider adding an IP address if you want to enable tunnel monitoring.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—tunnel.11<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—172.19.9.2/24<br>The configuration for VPN Peer B is:<br>• **Interface**—tunnel.12<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—192.168.69.2/24 |
| Step 3 Configure a static route, on the virtual router, to the destination subnet.<br> | 1. Select **Network > Virtual Router** and click the router you defined in step 4 above.<br>2. Select **Static Route**, click **Add**, and enter a new route to access the subnet that is at the other end of the tunnel.<br>In this example, the configuration for VPN Peer A is:<br>• **Destination**—192.168.69.0/24<br>• **Interface**—tunnel.11<br>The configuration for VPN Peer B is:<br>• **Destination**—172.19.9.0/24<br>• **Interface**—tunnel.12 |


© Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static Routing | |
|---|---|
| Step 4 Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).<br><br>Complete this task on both peers and make sure to set identical values. | 1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.<br> <br>2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.<br> |
| Step 5 Set up the IKE Gateway.<br> | 1. Select **Network > Network Profiles > IKE Gateway.**<br>2. Click **Add** and configure the options in the **General** tab**.**<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Local IP address**—192.168.210.26/24<br>• **Peer IP type/address**—static/192.168.210.120<br>• **Preshared keys**—enter a value<br>• **Local identification**—None; this means that the local IP address will be used as the local identification value.<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Local IP address**—192.168.210.120/24<br>• **Peer IP type/address**—static/192.168.210.26<br>• **Preshared keys**—enter same value as on Peer A<br>• **Local identification**—None<br>3. Select **Advanced Phase 1 Options** and select the IKE Crypto profile you created earlier to use for IKE phase 1. |

| Quick Config: Site-to-Site VPN with Static Routing | | |
|---|---|---|
| Step 6 | Set up the IPSec Tunnel. | 1. Select **Network > IPSec Tunnels**.<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Tunnel Interface**—tunnel.11<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IPSec Crypto profile defined in Step 4.<br>The configuration for VPN Peer B is:<br>• **Tunnel Interface**—tunnel.12<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IPSec Crypto defined in Step 4.<br>3. (Optional) Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity. Typically, the tunnel interface IP address for the VPN Peer is used.<br>4. (Optional) To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile. |
| Step 7 | Create policies to allow traffic between the sites (subnets). | 1. Select **Policies > Security**.<br>2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses. |
| Step 8 | Save any pending configuration changes. | Click **Commit**. |
| Step 9 | Test VPN connectivity. | See View the Status of the Tunnels. |

## Site-to-Site VPN with OSPF

In this example, each site uses OSPF for dynamic routing of traffic. The tunnel IP address on each VPN peer is statically assigned and serves as the next hop for routing traffic between the two sites.

 © Palo Alto Networks, Inc.

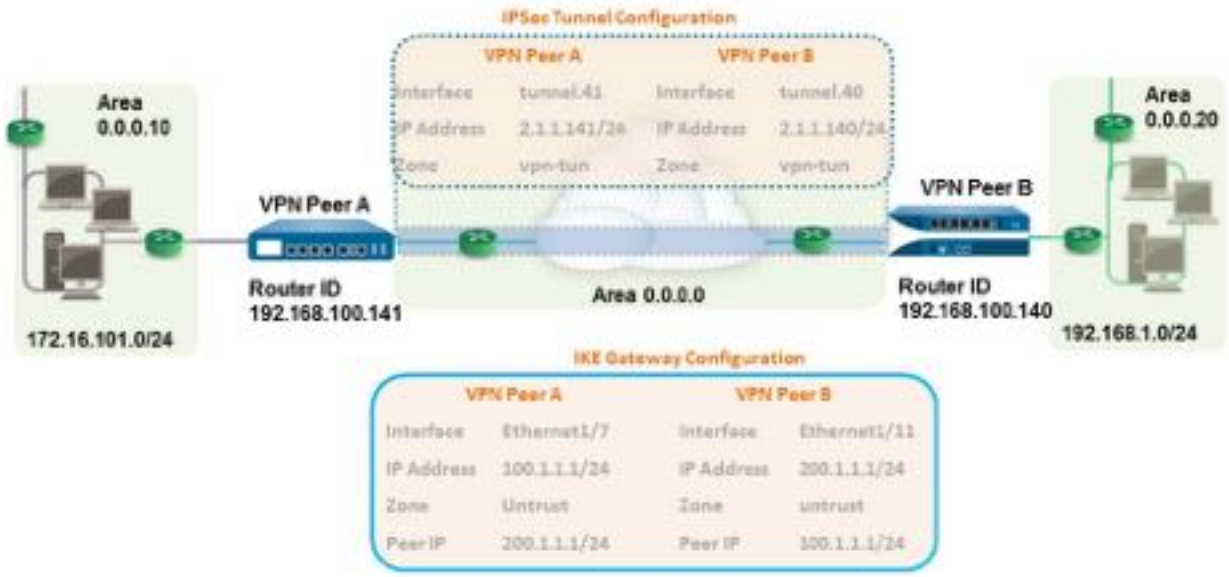


**Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF**

| | |
|---|---|
| Step 1 Configure the Layer 3 interfaces on each firewall. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—100.1.1.1/24<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—200.1.1.1/24 |

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | |
|---|---|
| Step 2 Create a tunnel interface and attach it to a virtual router and security zone.<br> | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, say, **.11**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.<br>This IP address will be used as the next hop IP address to route traffic to the tunnel and can also be used to monitor the status of the tunnel.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—tunnel.41<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.141/24<br>The configuration for VPN Peer B is:<br>• **Interface**—tunnel.40<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.140/24 |
| Step 3 Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).<br>Complete this task on both peers and make sure to set identical values. | 1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.<br> <br>2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.<br> |


© Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | | |
|---|---|---|
| Step 4 | Set up the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall.<br><br>For more information on the OSPF options that are available on the firewall, see Configure OSPF.<br><br>Use Broadcast as the link type when there are more than two OSPF routers that need to exchange routing information. | 1. Select **Network > Virtual Routers**, and select the default router or add a new router.<br>2. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.<br>3. In this example, the OSPF configuration for VPN Peer A is:<br>– **Router ID**: 192.168.100.141<br>– **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p<br>– **Area ID**: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast<br>The OSPF configuration for VPN Peer B is:<br>– **Router ID**: 192.168.100.140<br>– **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p<br>– **Area ID**: 0.0.0.20 that is assigned to the interface Ethernet1/15 and Link Type: Broadcast |
| Step 5 | Set up the IKE Gateway.<br><br>This examples uses static IP addresses for both VPN peers. Typically, the corporate office uses a statically configured IP address, and the branch side can be a dynamic IP address; dynamic IP addresses are not best suited for configuring stable services such as VPN. | 1. Select **Network > Network Profiles > IKE Gateway.**<br>2. Click **Add** and configure the options in the **General** tab**.**<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Local IP address**—100.1.1.1/24<br>• **Peer IP address**—200.1.1.1/24<br>• **Preshared keys**—enter a value<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Local IP address**—200.1.1.1/24<br>• **Peer IP address**—100.1.1.1/24<br>• **Preshared keys**—enter same value as on Peer A<br>3. Select the IKE Crypto profile you created earlier to use for IKE phase 1. |

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | |
|---|---|
| Step 6 Set up the IPSec Tunnel. | 1. Select **Network > IPSec Tunnels**.<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Tunnel Interface**—tunnel.41<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>The configuration for VPN Peer B is:<br>• **Tunnel Interface**—tunnel.40<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.<br>4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile. |
| Step 7 Create policies to allow traffic between the sites (subnets). | 1. Select **Policies > Security**.<br>2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses. |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | | |
|---|---|---|
| Step 8 | Verify OSPF adjacencies and routes from the CLI. | Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.<br>• **show routing protocol ospf neighbor**<br> 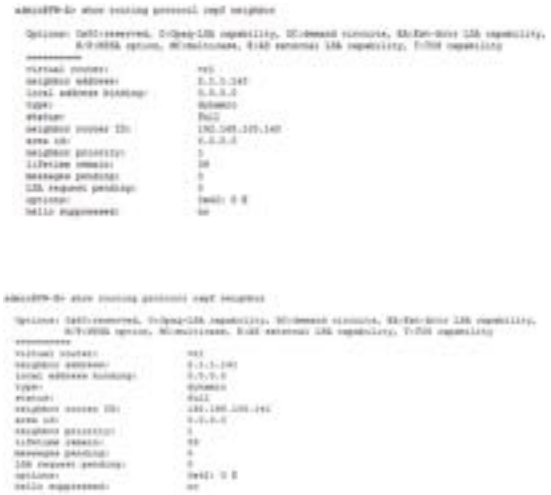 |
| | | • **show routing route type ospf**<br> 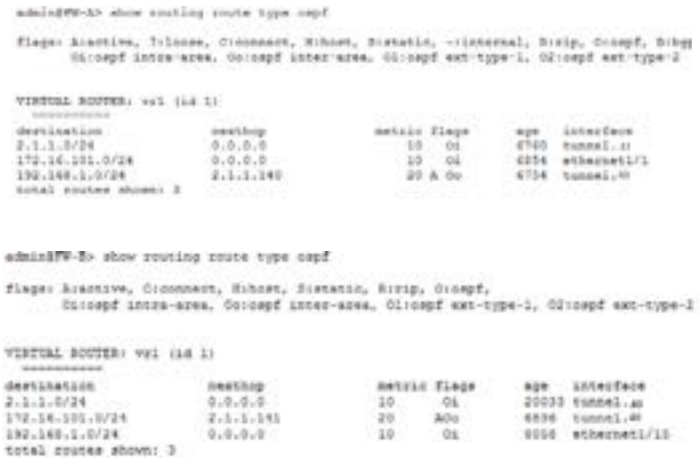 |
| Step 9 | Test VPN connectivity. | See Set up Tunnel Monitoring and View the Status of the Tunnels. |

## Site-to-Site VPN with Static and Dynamic Routing

In this example, one site uses static routes and the other site uses OSPF. When the routing protocol is not the same between the locations, the tunnel interface on each firewall must be configured with a static IP address. Then, to allow the exchange of routing information, the firewall that participates in both the static and dynamic routing process must be configured with a *Redistribution profile*. Configuring the redistribution profile enables the virtual router to redistribute and filter routes between protocols—static routes, connected routes, and hosts—

from the static autonomous system to the OSPF autonomous system. Without this redistribution profile, each protocol functions on its own and does not exchange any route information with other protocols running on the same virtual router.

In this example, the satellite office has static routes and all traffic destined to the 192.168.x.x network is routed to tunnel.41. The virtual router on VPN Peer B participates in both the static and the dynamic routing process and is configured with a redistribution profile in order to propagate (*export*) the static routes to the OSPF autonomous system.

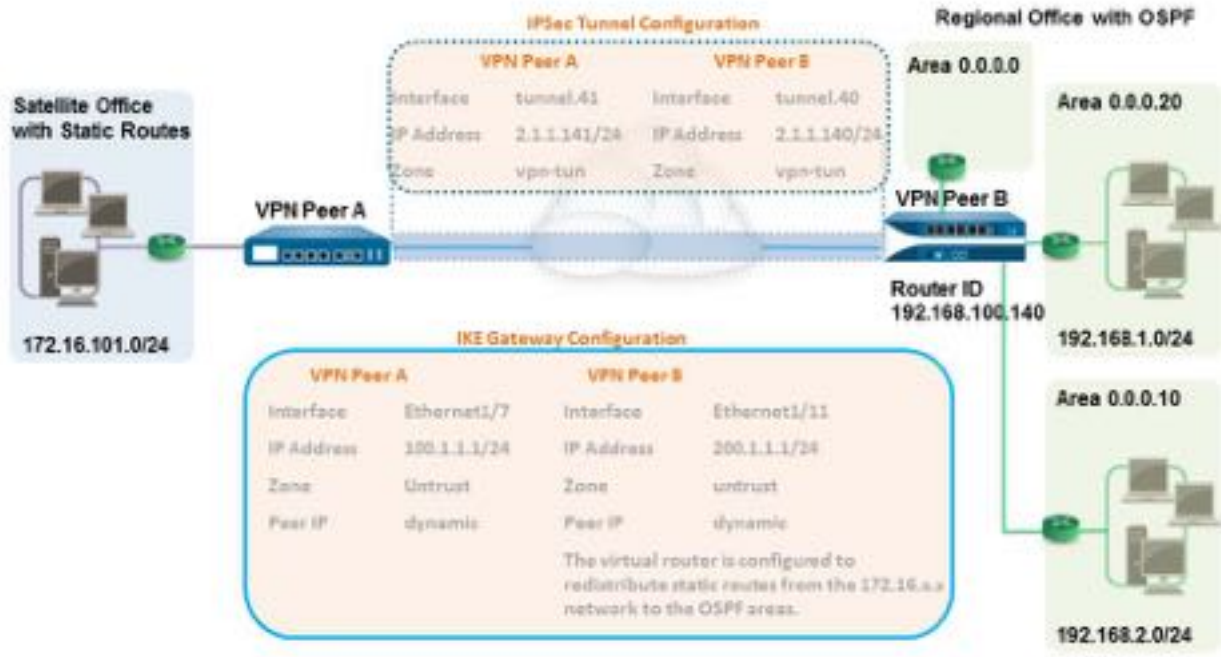


 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | | |
|---|---|---|
| Step 1 | Configure the Layer 3 interfaces on each firewall. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—100.1.1.1/24<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—200.1.1.1/24 |
| Step 2 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).<br>Complete this task on both peers and make sure to set identical values. | 1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.<br> <br>2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.<br> |

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 3 Set up the IKE Gateway.<br><br>With pre-shared keys, to add authentication scrutiny when setting up the IKE phase-1 tunnel, you can set up Local and Peer Identification attributes and a corresponding value that is matched in the IKE negotiation process. | 1. Select **Network > Network Profiles > IKE Gateway.**<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Local IP address**—100.1.1.1/24<br>• **Peer IP type**—dynamic<br>• **Preshared keys**—enter a value<br>• **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer A.<br>• **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer B<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Local IP address**—200.1.1.1/24<br>• **Peer IP address**—dynamic<br>• **Preshared keys**—enter same value as on Peer A<br>• **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer B<br>• **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer A<br>3. Select the IKE Crypto profile you created earlier to use for IKE phase 1. |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | | |
|---|---|---|
| Step 4 | Create a tunnel interface and attach it to a virtual router and security zone.<br> | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, say, **.41**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.<br>This IP address will be used to route traffic to the tunnel and to monitor the status of the tunnel.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—tunnel.41<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.141/24<br>The configuration for VPN Peer B is:<br>• **Interface**—tunnel.42<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.140/24 |
| Step 5 | Specify the interface to route traffic to a destination on the 192.168.x.x network. | 1. On VPN Peer A, select the virtual router.<br>2. Select **Static Routes**, and **Add** tunnel.41 as the **Interface** for routing traffic with a **Destination** in the 192.168.x.x network. |

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 6 Set up the static route and the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall. | 1. On VPN Peer B, select **Network > Virtual Routers**, and select the default router or add a new router.<br>2. Select **Static Routes** and **Add** the tunnel IP address as the next hop for traffic in the 172.168.x.x. network.<br>Assign the desired route metric; using a lower the value makes the a higher priority for route selection in the forwarding table.<br>3. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.<br>4. In this example, the OSPF configuration for VPN Peer B is:<br>• Router ID: 192.168.100.140<br>• Area ID: 0.0.0.0 is assigned to the interface Ethernet 1/12 Link type: Broadcast<br>• Area ID: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast<br>• Area ID: 0.0.0.20 is assigned to the interface Ethernet1/15 and Link Type: Broadcast |
| Step 7 Create a redistribution profile to inject the static routes into the OSPF autonomous system.<br> <br> | 1. Create a redistribution profile on VPN Peer B.<br>a. Select **Network > Virtual Routers**, and select the router you used above.<br>b. Select **Redistribution Profiles,** and click **Add**.<br>c. Enter a Name for the profile and select **Redist** and assign a **Priority** value. If you have configured multiple profiles, the profile with the lowest priority value is matched first.<br>d. Set **Source Type** as **static**, and click **OK**. The static route defined in Step 6-2 will be used for the redistribution.<br>2. Inject the static routes in to the OSPF system.<br>a. Select **OSPF> Export Rules** (for IPv4) or **OSPFv3> Export Rules** (for IPv6).<br>b. Click **Add**, and select the redistribution profile that you just created.<br>c. Select how the external routes are brought into the OSPF system. The default option, **Ext2** calculates the total cost of the route using only the external metrics. To use both internal and external OSPF metrics, use **Ext1.**<br>d. Assign a **Metric** (cost value) for the routes injected into the OSPF system. This option allows you to change the metric for the injected route as it comes into the OSPF system.<br>e. Click **OK** to save the changes. |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | | |
|---|---|---|
| Step 8 | Set up the IPSec Tunnel. | 1. Select **Network > IPSec Tunnels**.<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Tunnel Interface**—tunnel.41<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>The configuration for VPN Peer B is:<br>• **Tunnel Interface**—tunnel.40<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.<br>4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile. |
| Step 9 | Create policies to allow traffic between the sites (subnets). | 1. Select **Policies > Security**.<br>2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses. |

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 10 Verify OSPF adjacencies and routes from the CLI. | Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.<br>• **`show routing protocol ospf neighbor`**<br> 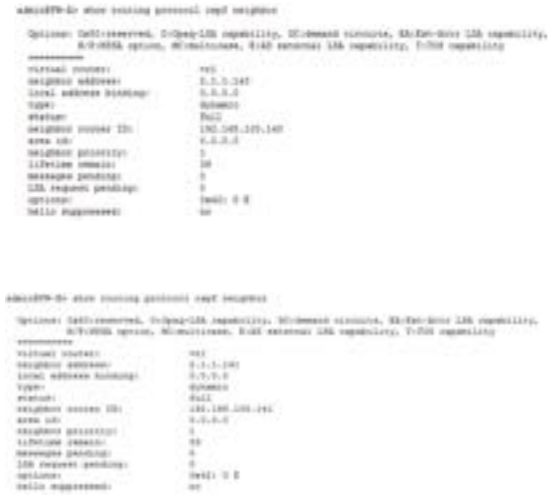 |
| | • **show routing route**<br>The following is an example of the output on each VPN peer.<br> |
| Step 11 Test VPN connectivity. | See Set up Tunnel Monitoring and View the Status of the Tunnels. |


© Palo Alto Networks, Inc.

# Large Scale VPN (LSVPN)

The GlobalProtect Large Scale VPN (LSVPN) feature on the Palo Alto Networks next-generation firewall simplifies the deployment of traditional hub and spoke VPNs, enabling you to quickly deploy enterprise networks with several branch offices with a minimum amount of configuration required on the remote *satellite* devices. This solution uses certificates for device authentication and IPSec to secure data.

LSVPN enables site-to-site VPNs between Palo Alto Networks firewalls. To set up a site-to-site VPN between a Palo Alto Networks firewall and another device, see VPNs.

The following topics describe the LSVPN components and how to set them up to enable site-to-site VPN services between Palo Alto Networks firewalls:

- LSVPN Overview
- Create Interfaces and Zones for the LSVPN
- Enable SSL Between GlobalProtect LSVPN Components
- Configure the Portal to Authenticate Satellites
- Configure GlobalProtect Gateways for LSVPN
- Configure the GlobalProtect Portal for LSVPN
- Prepare the Satellite Device to Join the LSVPN
- Verify the LSVPN Configuration
- LSVPN Quick Configs

# LSVPN Overview

GlobalProtect provides a complete infrastructure for managing secure access to corporate resources from your remote sites. This infrastructure includes the following components:

- **GlobalProtect Portal**—Provides the management functions for your GlobalProtect LSVPN infrastructure. Every satellite that participates in the GlobalProtect LSVPN receives configuration information from the portal, including configuration information to enable the satellites (the spokes) to connect to the gateways (the hubs). You configure the portal on an interface on any Palo Alto Networks next-generation firewall.
- **GlobalProtect Gateways**—A Palo Alto Networks firewall that provides the tunnel end point for satellite connections. The resources that the satellites access is protected by security policy on the gateway. It is not required to have a separate portal and gateway; a single firewall can function both as portal and gateway.
- **GlobalProtect Satellite**—A Palo Alto Networks firewall at a remote site that establishes IPSec tunnels with the gateway(s) at your corporate office(s) for secure access to centralized resources. Configuration on the satellite firewall is minimal, enabling you to quickly and easily scale your VPN as you add new sites.

The following diagram illustrates how the GlobalProtect LSVPN components work together.

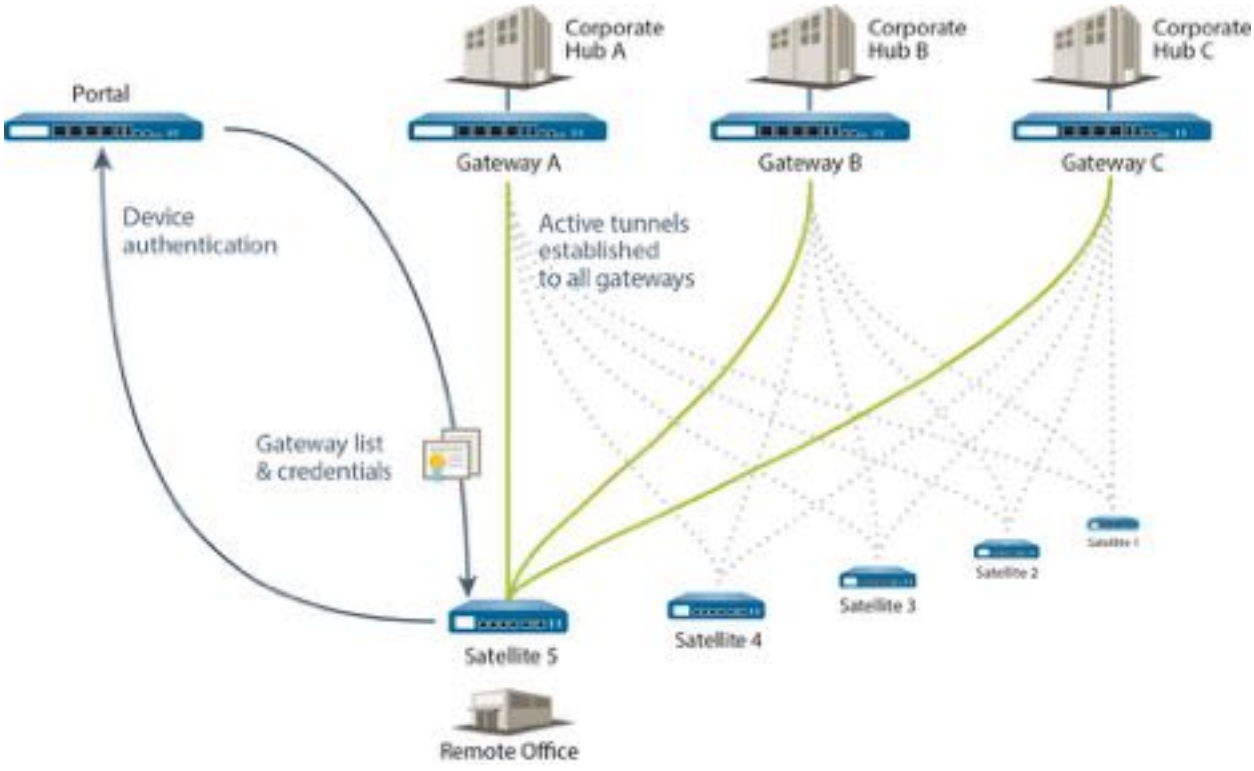


 © Palo Alto Networks, Inc.

# Create Interfaces and Zones for the LSVPN

You must configure the following interfaces and zones for your LSVPN infrastructure:

- **GlobalProtect portal**—Requires a Layer 3 interface for GlobalProtect satellites to connect to. If the portal and gateway are on the same firewall, they can use the same interface. The portal must be in a zone that is accessible from your branch offices.
- **GlobalProtect gateways**—Requires three interfaces: a Layer 3 interface in the zone that is reachable by the remote satellites, an internal interface in the trust zone that connects to the protected resources, and a logical tunnel interface for terminating the VPN tunnels from the satellites. Unlike other site-to-site VPN solutions, the GlobalProtect gateway only requires a single tunnel interface, which it will use for tunnel connections with all of your remote satellites (point-to-multipoint). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface.
- **GlobalProtect satellites**—Requires a single tunnel interface for establishing a VPN with the remote gateways (up to a maximum of 25 gateways). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface.

For more information about portals, gateways, and satellites see LSVPN Overview.

| Set Up Interfaces and Zones for the GlobalProtect LSVPN | |
|---|---|
| Step 1 Configure a Layer 3 interface.<br><br>The portal and each gateway and satellite all require a Layer 3 interface to enable traffic to be routed between sites.<br><br>If the gateway and portal are on the same firewall, you can use a single interface for both components.<br><br>IPv6 addresses are not supported with LSVPN. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for GlobalProtect LSVPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.11.100/24.<br>6. To save the interface configuration, click **OK**. |

| Set Up Interfaces and Zones for the GlobalProtect LSVPN (Continued) | | |
|---|---|---|
| Step 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.<br><br>IP addresses are not required on the tunnel interface unless plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.<br><br>Make sure to enable User-ID in the zone where the VPN tunnels terminate. | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *lsvpn-tun*), select the **Enable User Identification** check box, and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. (Optional) If you want to assign an IP address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.11.33/24.<br>6. To save the interface configuration, click **OK**. |
| Step 3 | If you created a separate zone for tunnel termination of VPN connections, create a security policy to enable traffic flow between the VPN zone and your trust zone. For example, the following policy rule enables traffic between the *lsvpn-tun* zone and the *L3-Trust* zone.<br> 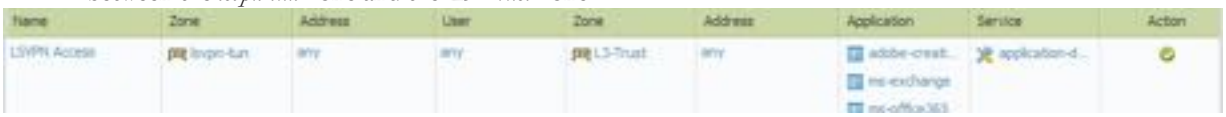 <br> | |
| Step 4 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# Enable SSL Between GlobalProtect LSVPN Components

All interaction between the GlobalProtect components occurs over an SSL connection. Therefore, you must generate and/or install the required certificates before configuring each component so that you can reference the appropriate certificate(s) and/or certificate profiles in the configurations for each component. The following sections describe the supported methods of certificate deployment, descriptions and best practice guidelines for the various GlobalProtect certificates, and provide instructions for generating and deploying the required certificates:

▲ About Certificate Deployment

▲ Deploy Server Certificates to the GlobalProtect LSVPN Components

## About Certificate Deployment

There are two basic approaches to deploying certificates for GlobalProtect LSVPN:

- **Enterprise Certificate Authority**—If you already have your own enterprise certificate authority, you can use this internal CA to issue an intermediate CA certificate for the GlobalProtect portal to enable it to issue certificates to the GlobalProtect gateways and satellites.
- **Self-Signed Certificates**—You can generate a self-signed root CA certificate on the firewall and use it to issue server certificates for the portal, gateway(s), and satellite(s). As a best practice, create a self-signed root CA certificate on the portal and use it to issue server certificates for the gateways and satellites. This way, the private key used for certificate signing stays on the portal.

## Deploy Server Certificates to the GlobalProtect LSVPN Components

The GlobalProtect LSVPN components use SSL/TLS to mutually authenticate. Before deploying the LSVPN, you must issue server certificates to the portal and gateways. You do not need to create server certificates for the satellite devices because the portal will issue a server certificate for each satellite during the first connection.

In addition, you must import the root CA certificate used to issue the server certificates onto each firewall that you plan to host as a gateway or satellite. Finally, on each gateway and satellite participating in the LSVPN, you must configure a certificate profile that will enable them to establish an SSL/TLS connection using mutual authentication.

The following workflow shows the best practice steps for deploying SSL certificates to the GlobalProtect LSVPN components:

You do not need to issue server certificates for the satellite devices because the portal will issue them as part of the satellite registration process.

| Deploy SSL Server Certificates to the GlobalProtect Components | | |
|---|---|---|
| Step 1 | On the firewall hosting the portal, create the root CA certificate for issuing self-signed certificates for the GlobalProtect components. | To use self-signed certificates, you must first create the root CA certificate that will be used to sign the certificates for the GlobalProtect components as follows:<br>1. To create a root CA certificate, select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**.<br>2. Enter a **Certificate Name**, such as *LSVPN_CA*. The certificate name cannot contain any spaces.<br>3. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed).<br>4. Select the **Certificate Authority** check box and then click **OK** to generate the certificate. |
| Step 2 | Generate server certificates for the GlobalProtect portal and gateway(s).<br>You must issue a unique self-signed server certificate for the portal and for each GlobalProtect gateway. The best practice is to issue all of the required certificates on the portal, so that the signing certificate (with the private key) does not have to be exported.<br> If the GlobalProtect portal and gateway are on the same firewall interface, you can use the same server certificate for both components. | Use the root CA on the portal to generate server certificates for each gateway you plan to deploy:<br>1. Select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**.<br>2. Enter a **Certificate Name**. The Certificate Name cannot contain any spaces.<br>3. Enter the FQDN (recommended) or IP address of the interface where you plan to configure the gateway in the **Common Name** field.<br>4. In the **Signed By** field, select the *LSVPN_CA* you created previously.<br>5. In the Certificate Attributes section, click **Add** and define the attributes to uniquely identify the gateway. Keep in mind that if you add a **Host Name** attribute (which populates the SAN field of the certificate), it must exactly match the value you defined for the **Common Name**.<br>6. **Generate** the certificate. |

 © Palo Alto Networks, Inc.

| Deploy SSL Server Certificates to the GlobalProtect Components (Continued) | |
|---|---|
| Step 3 Deploy the self-signed server certificates to the gateways.<br><br>**Best Practices:**<br>• Export the self-signed server certificates issued by the root CA from the portal and import them onto the gateways.<br>• Be sure to issue a unique server certificate for each gateway.<br>• The Common Name (CN) and, if applicable, the Subject Alternative Name (SAN) fields of the certificate must match the IP address or fully qualified domain name (FQDN) of the interface where you configure the gateway. | 1. On the portal, select **Device > Certificate Management > Certificates > Device Certificates**, select the gateway certificate you want to deploy, and click **Export**.<br>2. Select **Encrypted Private Key and Certificate (PKCS12)** from the **File Format** drop-down.<br>3. Enter (and re-enter) a **Passphrase** to encrypt the private key associated with the certificate and then click **OK** to download the PKCS12 file to your computer.<br>4. On the gateway, select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>5. Enter a **Certificate Name**.<br>6. Enter the path and name to the **Certificate File** you just downloaded from the portal, or **Browse** to find the file.<br>7. Select **Encrypted Private Key and Certificate (PKCS12)** as the **File Format**.<br>8. Enter the path and name to the PKCS12 file in the **Key File** field or **Browse** to find it.<br>9. Enter and re-enter the **Passphrase** you used to encrypt the private key when you exported it from the portal and then click **OK** to import the certificate and key. |
| Step 4 Import the root CA certificate used to issue server certificates for the LSVPN components.<br><br>You must import the root CA certificate onto all gateways and satellites. For security reasons, make sure you export the certificate only, and not the associated private key. | 1. Download the root CA certificate from the portal.<br>a. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>b. Select the root CA certificate used to issue certificates for the LSVPN components and click **Export**.<br>c. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (Do not export the private key.)<br>2. On the firewalls hosting the gateways and satellites, import the root CA certificate.<br>a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.<br>c. **Browse** to the **Certificate File** you downloaded from the CA.<br>d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.<br>e. Select the certificate you just imported on the **Device Certificates** tab to open it.<br>f. Select **Trusted Root CA** and then click **OK**.<br>g. **Commit** the changes. |

| Deploy SSL Server Certificates to the GlobalProtect Components (Continued) | | |
|---|---|---|
| Step 5 | Create a certificate profile.<br>The GlobalProtect LSVPN portal and each gateway require a certificate profile that specifies which certificate to use to authenticate the satellites. | 1. Select **Device > Certificate Management > Certificate Profile** and click **Add** and enter a profile **Name**.<br>2. Make sure **Username Field** is set to **None**.<br>3. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you imported in Step 4.<br>4. (Optional, but recommended) Enable use of CRL and/or OCSP to enable certificate status verification.<br>5. Click **OK** to save the profile. |
| Step 6 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

# Configure the Portal to Authenticate Satellites

In order to register with the LSVPN, each satellite must establish an SSL/TLS connection with the portal. After establishing the connection, the portal authenticates the satellite device to ensure that is authorized to join the LSVPN. After successfully authenticating the satellite, the portal will issue a server certificate for the satellite and push the LSVPN configuration specifying the gateways to which the satellite can connect and the root CA certificate required to establish an SSL connection with the gateways.

There are two ways that the satellite can authenticate to the portal during its initial connection:

- **Serial number**—You can configure the portal with the serial number of the satellite firewalls that are authorized to join the LSVPN. During the initial satellite connection to the portal, the satellite presents its serial number to the portal and if the portal has the serial number in its configuration, the satellite will be successfully authenticated. You add the serial numbers of authorized satellites when you configure the portal. See Configure the Portal.
- **Username and password**—If you would rather provision your satellites without manually entering the serial numbers of the satellite devices into the portal configuration, you can instead require the satellite administrator to authenticate when establishing the initial connection to the portal. Although the portal will always look for the serial number in the initial request from the satellite, if it cannot identify the serial number, the satellite administrator must provide a username and password to authenticate to the portal. Because the portal will always fall back to this form of authentication, you must create an authentication profile in order to commit the portal configuration. This requires that you set up an authentication profile for the portal LSVPN configuration even if you plan to authenticate satellites using the serial number.

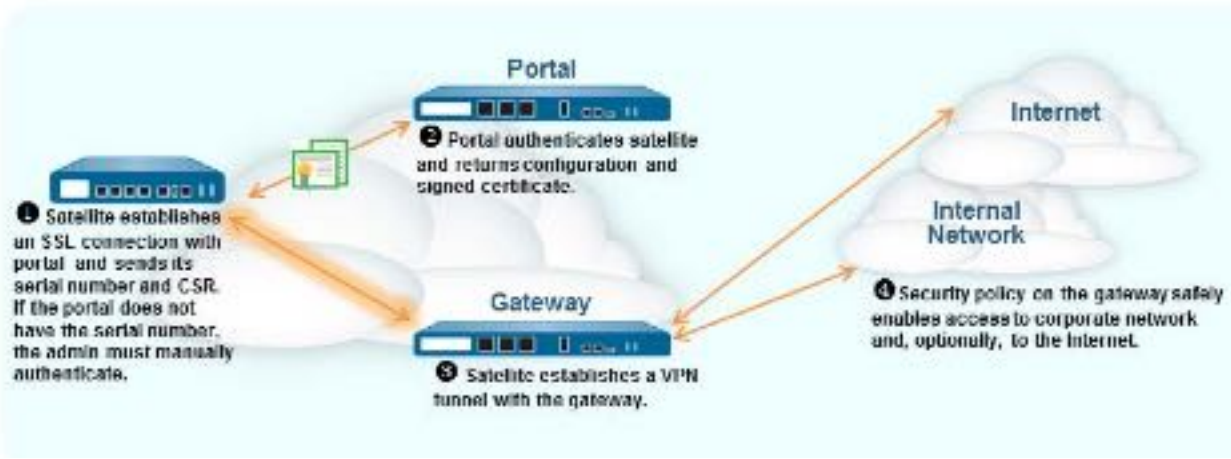


The following workflow describes how to set up the portal to authenticate satellites against an existing authentication service. GlobalProtect LSVPN supports external authentication using a local database, LDAP (including Active Directory), Kerberos, or RADIUS.

| Set Up Satellite Authentication | | |
|---|---|---|
| Step 1 | Create a server profile on the portal.<br><br>The server profile instructs the firewall how to connect to an external authentication service to validate the authentication credentials provided by the satellite administrator.<br><br>If you are using local authentication, you can skip this step and instead add a local user configuration for authenticating the satellite administrator.<br><br> If you are using LDAP to connect to Active Directory (AD), you must create a separate LDAP server profile for every AD domain. | 1. Select **Device > Server Profiles** and select type of profile (**LDAP**, **Kerberos**, or **RADIUS**).<br>2. Click **Add** and enter a **Name** for the profile, such as *LSVPN-Auth*.<br>3. (LDAP only) Select the **Type** of LDAP server you are connecting to.<br>4. Click **Add** in the Servers section and then enter information required to connect to the authentication service, including the server **Name**, **IP Address** (or FQDN), and **Port**.<br>5. (RADIUS and LDAP only) Specify settings to enable the firewall to authenticate to the authentication service as follows:<br>• RADIUS—Enter the shared **Secret** when adding the server entry.<br>• LDAP—Enter the **Bind DN** and **Bind Password**.<br>6. (LDAP and Kerberos only) Specify where to search for credentials in the directory service:<br>• LDAP—The **Base** DN specifies where in the LDAP tree to begin searching for users and groups. This field should populate automatically when you enter the server address and port. If it doesn't, check the service route to the LDAP server.<br>• Kerberos—Enter the Kerberos **Realm** name.<br>7. Specify the **Domain** name (without dots, for example `acme` not `acme.com`).<br>8. Click **OK** to save the server profile. |
| Step 2 | Create an authentication profile.<br><br>The authentication profile specifies which server profile to use to authenticate satellites. | 1. Select **Device > Authentication Profile** and click **Add**.<br>2. Enter a **Name** for the profile and then select the **Authentication** type (**Local Database, LDAP**, **Kerberos**, or **RADIUS**).<br>3. Select the **Server Profile** you created in Step 1.<br>4. (LDAP AD) Enter **sAMAccountName** as the **Login Attribute**.<br>5. Click **OK**. |
| Step 3 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# Configure GlobalProtect Gateways for LSVPN

Because the GlobalProtect configuration that the portal delivers to the satellites includes the list of gateways the satellite can connect to, it is a good idea to configure the gateways before configuring the portal.


- ▲ Prerequisite Tasks
- ▲ Configure the Gateway


## Prerequisite Tasks

Before you can configure the GlobalProtect gateway, you must have completed the following tasks:

- ❒ Created the interfaces (and zones) for the interface where you plan to configure each gateway. You must configure both the physical interface and the virtual tunnel interface. See Create Interfaces and Zones for the LSVPN.
- ❒ Set up the gateway server certificates and certificate profile required for enable GlobalProtect satellite and gateway to establish a mutual SSL/TLS connection. See Enable SSL Between GlobalProtect LSVPN Components.

## Configure the Gateway

After you have completed the Prerequisite Tasks, configure each GlobalProtect gateway to participate in the LSVPN as follows:

| Configure the Gateway for LSVPN | | |
|---|---|---|
| Step 1 | Add a gateway. | 1. Select **Network > GlobalProtect > Gateways** and click **Add**.<br>2. On the **General** tab, enter a **Name** for the gateway. The gateway name should not contain any spaces and as a best practice it should include the location or other descriptive information that will help identify the gateway.<br>3. (Optional) Select the virtual system to which this gateway belongs from the **Location** field. |
| Step 2 | Specify the network information to enable satellites to connect to the gateway.<br><br>If you have not yet created the network interface for the gateway, see Create Interfaces and Zones for the LSVPN for instructions. If you haven't yet created a server certificate for the gateway, see Deploy Server Certificates to the GlobalProtect LSVPN Components. | 1. Select the **Interface** that satellites will use for ingress access to the gateway.<br>2. Select the **IP Address** for gateway access.<br>3. Select the **Server Certificate** for the gateway from the drop-down. |

| Configure the Gateway for LSVPN (Continued) | | |
|---|---|---|
| Step 3 | Select the certificate profile for the gateway to use to authenticate satellites attempting to establish tunnels.<br><br>If you have not yet set up the certificate profile, see Enable SSL Between GlobalProtect LSVPN Components for instructions. | Select the **Certificate Profile** to you created for SSL communication between the LSVPN components. |
| Step 4 | Configure the tunnel parameters and enable tunneling. | 1. On the GlobalProtect Gateway dialog, select **Satellite Configuration > Tunnel Settings**.<br>2. Select the **Tunnel Configuration** check box to enable tunneling.<br>3. Select the **Tunnel Interface** you defined in Step 2 in Create Interfaces and Zones for the LSVPN.<br>4. (Optional) If you want to preserve the Type of Service (ToS) information in the encapsulated packets, select the **Copy TOS** check box. |
| Step 5 | (Optional) Enable tunnel monitoring.<br><br>Tunnel monitoring enables satellite devices to monitor its gateway tunnel connection, allowing it to failover to a backup gateway if the connection fails. Failover to another gateway is the only type of tunnel monitoring profile supported with LSVPN. | 1. Select the **Tunnel Monitoring** check box.<br>2. Specify the **Destination IP** address the satellite devices should use to determine if the gateway is active. Alternatively, if you configured an IP address for the tunnel interface, you can leave this field blank and the tunnel monitor will instead use the tunnel interface to determine if the connection is active.<br>3. Select **Failover** from the **Tunnel Monitor Profile** drop-down (this is the only supported tunnel monitor profile for LSVPN). |
| Step 6 | Select the Crypto Profile to use when establishing tunnel connections.<br><br>The crypto profile specifies the type IPSec encryption and/or the authentication method for securing the data that will traverse the tunnel. Because both tunnel endpoints in an LSVPN are trusted firewalls within your organization, you can typically use the default profile, which uses ESP-DH group2-AES 128 with SHA-1 encryption.<br><br>However, if you require a different mix of encryption and authentication mechanisms, you can optionally create a custom IPSec crypto profile. | Select **default** from the **IPSec Crypto Profile** drop-down or, optionally, select **New IPSec Crypto Profile** to define a new profile. For details on the authentication and encryption options in the crypto profile, refer to the online help. |

 © Palo Alto Networks, Inc.

| Configure the Gateway for LSVPN (Continued) | | |
|---|---|---|
| Step 7 | Configure the network settings to assign the satellites during establishment of the IPSec tunnel.<br><br> You can also configure the satellite device to push the DNS settings to its local clients by configuring a DHCP server on the firewall hosting the satellite. In this configuration, the satellite will push DNS settings it learns from the gateway to the DHCP clients. | 1. On the GlobalProtect Gateway dialog, select **Satellite Configuration > Network Settings**.<br>2. (Optional) If clients local to the satellite device need to resolve FQDNs on the corporate network, configure the gateway to push DNS settings to the satellites in one of the following ways:<br>• Manually define the **Primary DNS**, **Secondary DNS**, and **DNS Suffix** settings to push to the satellites.<br>• If the gateway has an interface that is configured as a DHCP client, you can set the **Inheritance Source** to that interface and the GlobalProtect satellites will be assigned the same settings received by the DHCP client.<br>3. To specify the **IP Pool** of addresses to assign the tunnel interface on the satellite devices when the VPN is established, click **Add** and then specify the IP address range(s) to use.<br>4. To define what destination subnets to route through the tunnel click **Add** in the **Access Route** area and then enter the routes as follows:<br>• If you want to route all traffic from the satellites through the tunnel, leave this field blank. Note that in this case, all traffic except traffic destined for the local subnet will be tunneled to the gateway.<br>• To route only some traffic through the gateway (called *split tunneling*), specify the destination subnets that must be tunneled. In this case, the satellite will route traffic that is not destined for a specified access route using its own routing table. For example, you may choose to only tunnel traffic destined for your corporate network, and use the local satellite to safely enable Internet access.<br>• If you want to enable routing between satellites, enter the summary route for the network protected by each satellite. |
| Step 8 | (Optional) Define what routes, if any, the gateway will accept from satellites.<br><br>By default, the gateway will not add any routes satellites advertise to its routing table. If you do not want the gateway to accept routes from gateways, you do not need to complete this step. | 1. To enable the gateway to accept routes advertised by satellites, select **Satellite Configuration > Route Filter**.<br>2. Select the **Accept published routes** check box.<br>3. To filter which of the routes advertised by the satellites to add to the gateway routing table, click **Add** and then define the subnets to include. For example, if all the satellites are configured with subnet 192.168.x.0/24 on the LAN side, configuring a permitted route of 192.168.0.0/16 to enable the gateway to only accept routes from the satellite if it is in the 192.168.0.0/16 subnet. |
| Step 9 | Save the gateway configuration. | 1. Click **OK** to save the settings and close the GlobalProtect Gateway dialog.<br>2. **Commit** the configuration. |

# Configure the GlobalProtect Portal for LSVPN

The GlobalProtect portal provides the management functions for your GlobalProtect LSVPN. Every satellite system that participates in the LSVPN receives configuration information from the portal, including information about available gateways as well as the certificate it needs in order to connect to the gateways.

The following sections provide procedures for setting up the portal:

- Prerequisite Tasks
- Configure the Portal
- Define the Satellite Configurations

## Prerequisite Tasks

Before you can configure the GlobalProtect portal, you must have completed the following tasks:

- Created the interfaces (and zones) for the firewall interface where you plan to configure the portal. See Create Interfaces and Zones for the LSVPN.
- Issued the portal server certificate, gateway server certificates, and set up the portal to issue server certificates for the satellites. See Enable SSL Between GlobalProtect LSVPN Components.
- Defined the authentication profile that will be used to authenticate GlobalProtect satellites in the event that the serial number is not available. See Configure the Portal to Authenticate Satellites.
- Configured the global protect gateways. See Configure GlobalProtect Gateways for LSVPN.

 © Palo Alto Networks, Inc.

## Configure the Portal

After you have completed the Prerequisite Tasks, configure the GlobalProtect portal as follows:

| Configure the Portal for LSVPN | | |
|---|---|---|
| Step 1 | Add the portal. | 1. Select **Network > GlobalProtect > Portals** and click **Add**.<br>2. On the **Portal Configuration** tab, enter a **Name** for the portal. The portal name should not contain any spaces.<br>3. (Optional) Select the virtual system to which this portal belongs from the **Location** field. |
| Step 2 | Specify the network information to enable satellites to connect to the portal.<br><br>If you have not yet created the network interface for the portal, see Create Interfaces and Zones for the LSVPN for instructions. If you haven't yet created a server certificate for the portal and issued gateway certificates, see Deploy Server Certificates to the GlobalProtect LSVPN Components. | 1. Select the **Interface** that satellites will use for ingress access to the portal.<br>2. Select the **IP Address** for satellite access to the portal.<br>3. Select the **Server Certificate** you generated to enable the satellite to establish SSL connection with the portal. |
| Step 3 | Specify an authentication profile for authenticating satellite devices.<br><br>Even if you plan to manually configure the portal with the serial numbers of the satellites, you must define an authentication profile or you will not be able to save the configuration. | • Select the **Authentication Profile** you defined for authenticating satellites.<br>• If you have not yet set up the authentication profile, select **New Authentication Profile** to create one now. See Configure the Portal to Authenticate Satellites for instructions. If the portal is unable to validate the serial number of a connecting satellite, it will fall back to the authentication profile and therefore you must configure an authentication profile in order to save the portal configuration. |
| Step 4 | Continue with defining the configurations to push to the satellites, or, if you have already created the satellite configurations, save the portal configuration. | Click **OK** to save the portal configuration or continue to Define the Satellite Configurations. |

## Define the Satellite Configurations

When a GlobalProtect satellite connects and successfully authenticates to the GlobalProtect portal, the portal delivers a satellite configuration, which specifies what gateways the satellite can connect to. If all your satellites will use the same gateway and certificate configurations, you can create a single satellite configuration to deliver to all satellites upon successful authentication. However, if you require different satellite configurations—for example if you want one group of satellites to connect to one gateway and another group of satellites to connect to a different gateway—you can create a separate satellite configuration for each. The portal will then use the enrollment username/group name or the serial number of the satellite device to determine which satellite configuration to deploy. As with security rule evaluation, the portal looks for a match starting from the top of the list. When it finds a match, it delivers the corresponding configuration to the satellite.

For example, the following figure shows a network in which some branch offices require VPN access to the corporate applications protected by your perimeter firewalls and another site needs VPN access to the datacenter.

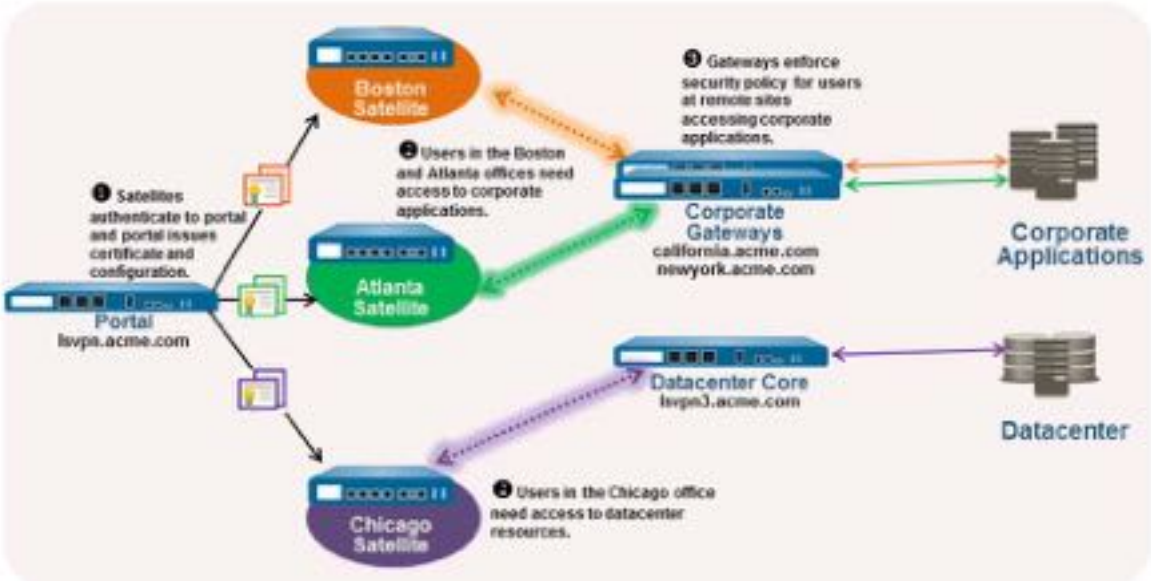


Use the following procedure to create one or more satellite configurations.

 © Palo Alto Networks, Inc.

| Create a GlobalProtect Satellite Configuration | | |
|---|---|---|
| Step 1 | Specify the certificates required to enable satellites to participate in the LSVPN. | 1. Select **Network > GlobalProtect > Portals** and select the portal configuration for which you want to add a satellite configuration and then select the **Satellite Configuration** tab.<br>2. In the **Trusted Root CA** field, click **Add** and then select the CA certificate used to issue the gateway server certificates. The portal will deploy the root CA certificate you add here to all satellites as part of the configuration to enable the satellite to establish an SSL connection with the gateways. As a best practice, all of your gateways should use the same issuer.<br> If the root CA certificate used to issue your gateway server certificates is not on the portal, you can **Import** it now. See Enable SSL Between GlobalProtect LSVPN Components for details on how to import a root CA certificate.<br>3. Select the Root CA certificate that the portal will use to issue certificates to satellites upon successfully authenticating them from the **Issuing Certificate** drop-down. |
| Step 2 | Add a satellite configuration.<br><br>The satellite configuration specifies the GlobalProtect LSVPN configuration settings to deploy to the connecting satellites. You must define at least one satellite configuration. | In the Satellite Configuration section, click **Add** and enter a **Name** for the configuration.<br><br>If you plan to create multiple configurations, make sure the name you define for each is descriptive enough to allow you to distinguish them. |
| Step 3 | Specify which satellites to deploy this configuration to. There are two ways to specify which satellites will get the configuration: by enrollment user/group name and/or using the serial number of the satellite devices.<br><br>The portal uses the **Enrollment User/User Group** settings and/or **Devices** serial numbers to match a satellite to a configuration. Therefore, if you have multiple configurations, be sure to order them properly. As soon as the portal finds a match, it will deliver the configuration. Therefore, more specific configurations must precede more general ones. See Step 6 for instructions on ordering the list of satellite configurations. | Specify the match criteria for the satellite configuration as follows:<br>• To restrict this configuration to satellite devices with specific serial numbers, select the **Devices** tab, click **Add**, and enter serial number (you do not need to enter the satellite hostname; it will be automatically added when the satellite connects). Repeat this step for each satellite you want to receive this configuration.<br>• Select the **Enrollment User/User Group** tab, click **Add**, and then select the user or group you want to receive this configuration. Satellites that do not match on serial number will be required to authenticate as a user specified here (either an individual user or group member).<br><br> Before you can restrict the configuration to specific groups, you must Map Users to Groups. |

| Create a GlobalProtect Satellite Configuration (Continued) | | |
|---|---|---|
| Step 4 | Specify the gateways that satellites with this configuration can establish VPN tunnels with.<br><br> Routes published by the gateway are installed on the satellite as static routes. The metric for the static route is 10x the routing priority. If you have more than one gateway, make sure to also set the routing priority to ensure that routes advertised by backup gateways have higher metrics compared to the same routes advertised by primary gateways. For example, if you set the routing priority for the primary gateway and backup gateway to 1 and 10 respectively, the satellite will use 10 as the metric for the primary gateway and 100 as the metric for the backup gateway. | 1. On the **Gateways** tab, click **Add**.<br>2. Enter a descriptive **Name** for the gateway. The name you enter here should match the name you defined when you configured the gateway and should be descriptive enough identify the location of the gateway.<br>3. Enter the FQDN or IP address of the interface where the gateway is configured in the **Gateways** field. The address you specify must exactly match the Common Name (CN) in the gateway server certificate.<br>4. (Optional) If you are adding two or more gateways to the configuration, the **Routing Priority** helps the satellite pick the preferred gateway. Enter a value in the range of 1-25, with lower numbers having the higher priority (that is, the gateway the satellite will connect to if all gateways are available). The satellite will multiply the routing priority by 10 to determine the routing metric. |
| Step 5 | Save the satellite configuration. | 1. Click **OK** to save the satellite configuration.<br>2. If you want to add another satellite configuration, repeat Step 2 through Step 5. |
| Step 6 | Arrange the satellite configurations so that the proper configuration is deployed to each satellite. | • To move a satellite configuration up on the list of configurations, select the configuration and click **Move Up**.<br>• To move a satellite configuration down on the list of configurations, select the configuration and click **Move Down**. |
| Step 7 | Save the portal configuration. | 1. Click **OK** to save the settings and close the GlobalProtect Portal dialog.<br>2. **Commit** your changes. |

 © Palo Alto Networks, Inc.

# Prepare the Satellite Device to Join the LSVPN

In order to participate in the LSVPN, the satellite devices require a minimal amount of configuration. Because the required configuration is minimal, you can pre-configure the devices before shipping them to your branch offices for installation.

| Prepare the Satellite Device to Join the GlobalProtect LSVPN | | |
|---|---|---|
| Step 1 | Configure a Layer 3 interface. | This is the physical interface the satellite will use to connect to the portal and the gateway. This interface must be in a zone that allows access outside of the local trust network. As a best practice, create a dedicated zone for VPN connections for visibility and control over traffic destined for the corporate gateways. |
| Step 2 | Configure the logical tunnel interface for the tunnel to use to establish VPN tunnels with the GlobalProtect gateways.<br> IP addresses are not required on the tunnel interface unless plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues. | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select an existing zone or create a separate zone for VPN tunnel traffic by clicking **New Zone** and defining a **Name** for new zone (for example *lsvpnsat*).<br>4. In the **Virtual Router** drop-down, select **default**.<br>5. (Optional) If you want to assign an IP address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 2.2.2.11/24.<br>6. To save the interface configuration, click **OK**. |

| Prepare the Satellite Device to Join the GlobalProtect LSVPN (Continued) | | |
|---|---|---|
| Step 3 | If you generated the portal server certificate using a Root CA that is not trusted by the satellites (for example, if you used self-signed certificates), import the root CA certificate used to issue the portal server certificate.<br><br>The root CA certificate is required to enable the satellite device to establish the initial connection with the portal to obtain the LSVPN configuration. | 1. Download the CA certificate that was used to generate the portal server certificates. If you are using self-signed certificates, export the root CA certificate from the portal as follows:<br>a. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>b. Select the CA certificate, and click **Export**.<br>c. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (You do not need to export the private key.)<br>2. Import the root CA certificate you just exported onto each satellite as follows.<br>a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.<br>c. **Browse** to the **Certificate File** you downloaded from the CA.<br>d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.<br>e. Select the certificate you just imported on the **Device Certificates** tab to open it.<br>f. Select **Trusted Root CA** and then click **OK**. |
| Step 4 | Configure the IPSec tunnel configuration. | 1. Select **Network > IPSec Tunnels** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the IPSec configuration.<br>3. Select the **Tunnel Interface** you created for the satellite.<br>4. Select **GlobalProtect Satellite** as the **Type**.<br>5. Enter the IP address or FQDN of the portal as the **Portal Address**.<br>6. Select the Layer 3 **Interface** you configured for the satellite.<br>7. Select the **Local IP Address** to use on the selected interface. |
| Step 5 | (Optional) Configure the satellite to publish local routes to the gateway.<br><br>Pushing routes to the gateway enables traffic to the subnets local to the satellite via the gateway. However, the gateway must also be configured to accept the routes. | 1. To enable the satellite to push routes to the gateway, on the **Advanced** tab select **Publish all static and connected routes to Gateway**.<br>2. (Optional) If you only want to push routes for specific subnets rather than all routes, click **Add** in the Subnet section and specify which subnet routes to publish. |
| Step 6 | Save the satellite configuration. | 1. Click **OK** to save the IPSec tunnel settings.<br>2. Click **Commit**. |

 © Palo Alto Networks, Inc.

| Prepare the Satellite Device to Join the GlobalProtect LSVPN (Continued) | |
|---|---|
| Step 7 If required, provide the credentials to allow the satellite to authenticate to the portal.<br><br>This step is only required if the portal was unable to find a serial number match in its configuration or if the serial number didn't work. In this case, the satellite will not be able to establish the tunnel with the gateway(s). | 1. Select **Network > IPSec Tunnels** and click the **Gateway Info** link in the Status column of the tunnel configuration you created for the LSVPN.<br>2. Click the **enter credentials** link in the **Portal Status** field and username and password required to authenticate the satellite to the portal.<br><br>After the portal successfully authenticates to the portal, it will receive its signed certificate and configuration, which it will use to connect to the gateway(s). You should see the tunnel establish and the **Status** change to **Active**.<br><br> |

# Verify the LSVPN Configuration

After configuring the portal, gateways, and satellite devices, verify that the satellites are able to connect to the portal and gateway and establish VPN tunnels with the gateway(s).

| Verify the LSVPN Configuration | | |
|---|---|---|
| Step 1 | Verify satellite connectivity with portal. | From the firewall hosting the portal, verify that satellites are successfully connecting by selecting **Network > GlobalProtect > Portal** and clicking **Satellite Info** in the Info column of the portal configuration entry.<br> |
| Step 1 | Verify satellite connectivity with the gateway(s). | On each firewall hosting a gateway, verify that satellites are able to establish VPN tunnels by selecting **Network > GlobalProtect > Gateways** and click **Satellite Info** in the Info column of the gateway configuration entry. Satellites that have successfully established tunnels with the gateway will display on the **Active Satellites** tab.<br> |
| Step 1 | Verify LSVPN tunnel status on the satellite. | On each firewall hosting a satellite, verify the tunnel status by selecting **Network > IPSec Tunnels** and verify active Status as indicated by a green icon.<br> |


© Palo Alto Networks, Inc.

# LSVPN Quick Configs

The following sections provide step-by-step instructions for configuring some common GlobalProtect LSVPN deployments:

- Basic LSVPN Configuration with Static Routing
- Advanced LSVPN Configuration with Dynamic Routing

# Basic LSVPN Configuration with Static Routing

This quick config shows the fastest way to get up and running with LSVPN. In this example, a single firewall at the corporate headquarters site is configured as both a portal and a gateway. Satellite devices can be quickly and easily deployed with minimal configuration for optimized scalability.

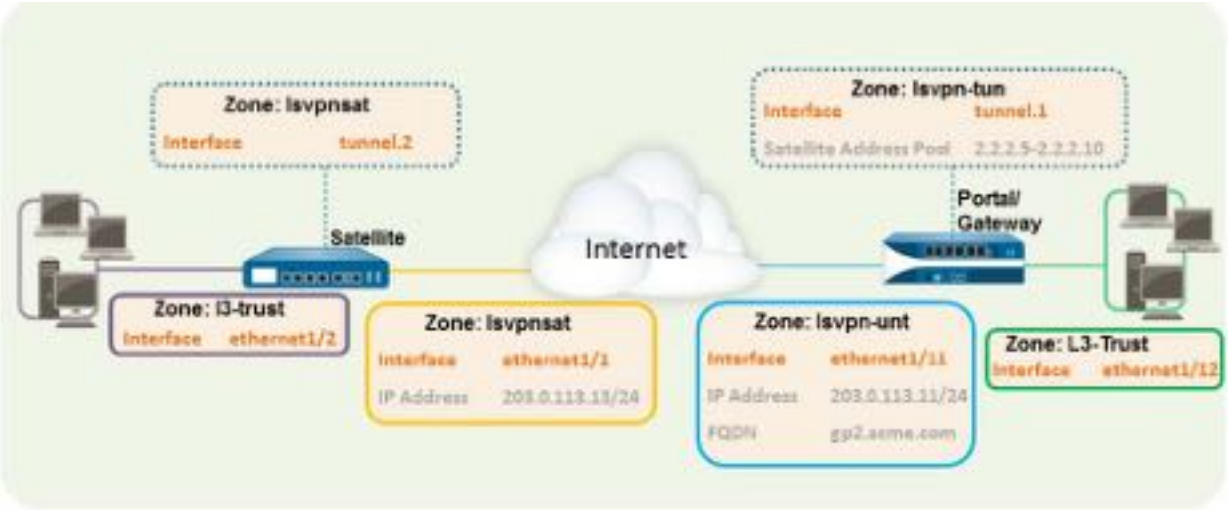


The following workflow shows the steps for setting up this basic configuration:

| Quick Config: Basic LSVPN with Static Routing | | |
|---|---|---|
| Step 1 | Configure a Layer 3 interface. | In this example, the Layer 3 interface on the portal/gateway requires the following configuration:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—lsvpn-unt<br>• **IPv4**—203.0.113.11/24 |
| Step 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.<br>To enable visibility into users and groups connecting over the VPN, enable User-ID in the zone where the VPN tunnels terminate. | In this example, the Tunnel interface on the portal/gateway requires the following configuration:<br>• **Interface**—tunnel.1<br>• **Security Zone**—lsvpn-tun |
| Step 3 | Create the security policy rule to enable traffic flow between the VPN zone where the tunnel terminates (lsvpn-tun) and the trust zone where the corporate applications reside (L3-Trust).<br> 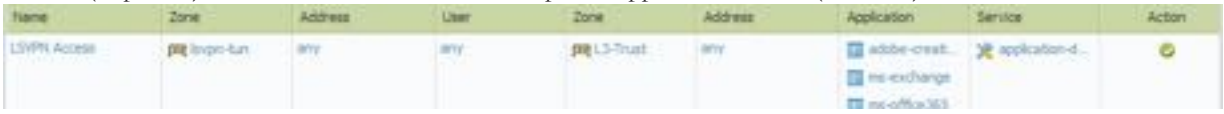 | |



 © Palo Alto Networks, Inc.

| Quick Config: Basic LSVPN with Static Routing (Continued) | | |
|---|---|---|
| Step 4 | Issue a self-signed server certificate for the portal/gateway.<br><br>The certificate subject name must match the FQDN or IP address of the Layer 3 interface you create for the portal/gateway. | 1. On the firewall hosting the portal, create the root CA certificate for issuing self-signed certificates for the GlobalProtect components. In this example, the root CA certificate, *lsvpn-CA*, will be used to issue the server certificate for the portal/gateway. In addition, the portal will use this root CA certificate to sign the CSRs from the satellite devices.<br>2. Generate server certificates for the GlobalProtect portal and gateway(s).<br><br>Because the portal and gateway will be on the same interface in this example, they can share a server certificate. In this example, the server certificate is named *lsvpnserver*. |
| Step 5 | Create a certificate profile. | In this example, the certificate profile, *lsvpn-profile*, references the root CA certificate *lsvpn-CA*. The gateway will use this certificate profile to authenticate satellites attempting to establish VPN tunnels. |
| Step 6 | Configure an authentication profile for the portal to use if the satellite serial number is not available. | 1. Create a server profile on the portal.<br>2. Create an authentication profile. In this example, the profile *lsvpn-sat* is used to authenticate satellites. |
| Step 7 | Configure the Gateway for LSVPN. | Select **Network > GlobalProtect > Gateways** and **Add** a configuration. This example requires the following gateway configuration:<br>• **Interface**—ethernet1/11<br>• **IP Address**—203.0.113.11/24<br>• **Server Certificate**—lsvpnserver<br>• **Certificate Profile**—lsvpn-profile<br>• **Tunnel Interface**—tunnel.1<br>• **Primary DNS**/**Secondary DNS**—4.2.2.1/4.2.2.2<br>• **IP Pool**—2.2.2.111-2.2.2.120<br>• **Access Route**—10.2.10.0/24 |
| Step 8 | Configure the Portal for LSVPN. | Select **Network > GlobalProtect > Portal** and **Add** a configuration. This example requires the following portal configuration:<br>• **Interface**—ethernet1/11<br>• **IP Address**—203.0.113.11/24<br>• **Server Certificate**—lsvpnserver<br>• **Authentication Profile**—lsvpn-sat |

| Quick Config: Basic LSVPN with Static Routing (Continued) | |
|---|---|
| Step 9 Create a GlobalProtect Satellite Configuration. | On the **Satellite Configuration** tab in the portal configuration, **Add** a Satellite Configuration and a Trusted Root CA and specify the CA the portal will use to issue certificates for the satellites. In this example the required settings are as following:<br>• **Gateway**—203.0.113.11<br>• **Issuing Certificate**—lsvpn-CA<br>• **Trusted Root CA**—lsvpn-CA |
| Step 10 Prepare the Satellite Device to Join the LSVPN. | The satellite configuration in this example requires the following settings:<br>Interface Configuration<br>• Layer 3 interface—ethernet1/1, 203.0.113.13/24<br>• Tunnel interface—tunnel.2<br>• Zone—lsvpnsat<br>Root CA Certificate from Portal<br>• lsvpn-CA<br>IPSec Tunnel Configuration<br>• **Tunnel Interface**—tunnel.2<br>• **Portal Address**—203.0.113.11<br>• **Interface**—ethernet1/1<br>• **Local IP Address**—203.0.113.13/24<br>• **Publish all static and connected routes to Gateway**—enabled |

 © Palo Alto Networks, Inc.

# Advanced LSVPN Configuration with Dynamic Routing

In larger LSVPN deployments with multiple gateways and many satellites, investing a little more time in the initial configuration to set up dynamic routing will simplify the maintenance of gateway configurations because access routes will update dynamically. The following example configuration shows how to extend the basic LSVPN configuration to configure OSPF as the dynamic routing protocol.

Setting up an LSVPN to use OSPF for dynamic routing requires the following additional steps on the gateways and the satellites:

- Manual assignment of IP addresses to tunnel interfaces on all gateways and satellites.
- Configuration of OSPF point-to-multipoint (P2MP) on the virtual router on all gateways and satellites. In addition, as part of the OSPF configuration on each gateway, you must manually define the tunnel IP address of each satellite as an OSPF neighbor. Similarly, on each satellite, you must manually define the tunnel IP address of each gateway as an OSPF neighbor.

Although dynamic routing requires additional setup during the initial configuration of the LSVPN, it reduces the maintenance tasks associated with keeping routes up to date as topology changes occur on your network.

The following figure shows an LSVPN dynamic routing configuration. This example shows how to configure OSPF as the dynamic routing protocol for the VPN.

For a basic setup of a LSVPN, follow the steps in Basic LSVPN Configuration with Static Routing. You can then complete the steps in the following workflow to extend the configuration to use dynamic routing rather than static routing.

| Quick Config: LSVPN with Dynamic Routing | |
|---|---|
| Step 1 Add an IP address to the tunnel interface configuration on each gateway and each satellite. | Complete the following steps on each gateway and each satellite:<br>1. Select **Network > Interfaces > Tunnel** and select the tunnel configuration you created for the LSVPN to open the Tunnel Interface dialog.<br>If you have not yet created the tunnel interface, see Step 2 in Quick Config: Basic LSVPN with Static Routing.<br>2. On the **IPv4** tab, click **Add** and then enter an IP address and subnet mask. For example, to add an IP address for the gateway tunnel interface you would enter 2.2.2.100/24.<br>3. Click **OK** to save the configuration. |
| Step 2 Configure the dynamic routing protocol on the gateway.<br> 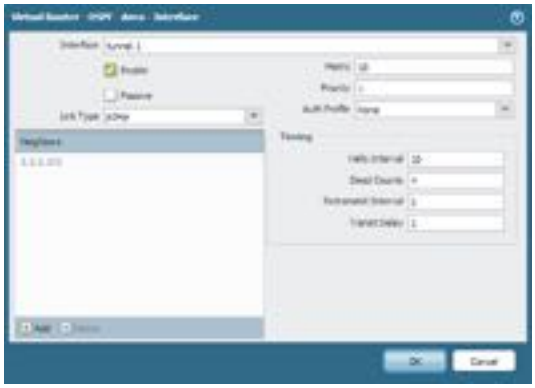 | To configure OSPF on the gateway:<br>1. Select **Network > Virtual Routers** and select the virtual router associated with your VPN interfaces.<br>2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.<br>3. If you are creating a new area, enter an **Area ID** on the **Type** tab.<br>4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.<br>5. Select **p2mp** as the **Link Type**.<br>6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each satellite device, for example 2.2.2.111.<br>7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.<br>8. Repeat this step each time you add a new satellite to the LSVPN. |
| Step 3 Configure the dynamic routing protocol on the satellite.<br> 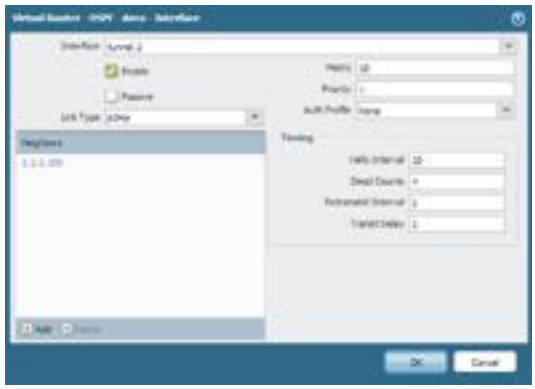 | To configure OSPF on the satellite:<br>1. Select **Network > Virtual Routers** and select the virtual router associated with your VPN interfaces.<br>2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.<br>3. If you are creating a new area, enter an **Area ID** on the **Type** tab.<br>4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.<br>5. Select **p2mp** as the **Link Type**.<br>6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each GlobalProtect gateway, for example 2.2.2.100.<br>7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.<br>8. Repeat this step each time you add a new gateway. |

 © Palo Alto Networks, Inc.

**Quick Config: LSVPN with Dynamic Routing (Continued)**

Step 4 Verify that the gateways and satellites are able to form router adjacencies.

- On each satellite and each gateway, confirm that peer adjacencies have formed and that routing table entries have been created for the peers (that is, the satellites have routes to the gateways and the gateways have routes to the satellites). Select **Network > Virtual Router** and click the **More Runtime Stats** link for the virtual router you are using for the LSVPN. On the Routing tab, verify that the LSVPN peer has a route.

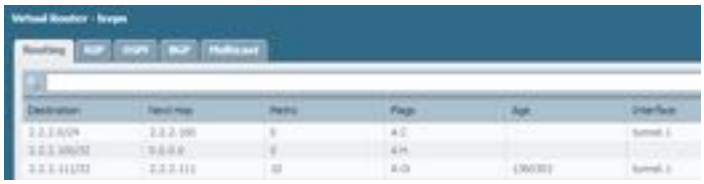

- On the **OSPF > Interface** tab, verify that the **Type** is **p2mp**.

- On the **OSPF > Neighbor** tab, verify that the firewalls hosting your gateways have established router adjacencies with the firewalls hosting your satellites and vice versa. Also verify that the **Status** is **Full**, indicating that full adjacencies have been established.

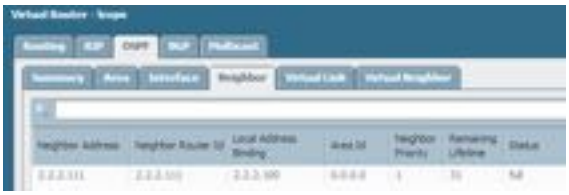

 © Palo Alto Networks, Inc.

# Networking

All Palo Alto Networks next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, enabling you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports. The Interface Deployments section provides basic information on each type of deployment. For more detailed deployment information, refer to Designing Networks with Palo Alto Networks Firewalls.

The following topics describe networking concepts and how to integrate Palo Alto Networks next-generation firewalls into your network.

- Interface Deployments
- Virtual Routers
- Static Routes
- RIP
- OSPF
- BGP
- Session Settings and Timeouts
- DHCP
- NAT
- LACP

For information on route distribution, refer to Understanding Route Redistribution and Filtering.

# Interface Deployments

A Palo Alto Networks firewall can operate in multiple deployments at once because the deployments occur at the interface level. The following sections describe the deployments.


- Virtual Wire Deployments
- Layer 2 Deployments
- Layer 3 Deployments
- Tap Mode Deployments


## Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together and should be used only when no switching or routing is needed.

A virtual wire deployment allows the following conveniences:

- Simplifies installation and configuration.
- Does not require any configuration changes to surrounding or adjacent network devices.

The "default-vwire" that is shipped as the factory default configuration binds together Ethernet ports 1 and 2 and allows all untagged traffic. You can, however, use a virtual wire to connect any two ports and configure it to block or allow traffic based on the virtual LAN (VLAN) tags; the VLAN tag "0" indicates untagged traffic. You can also create multiple subinterfaces, add them into different zones and then classify traffic according to a VLAN tag, or a combination of a VLAN tag with IP classifiers (address, range, or subnet) to apply granular policy control for specific VLAN tags or for VLAN tags from a specific source IP address, range, or subnet.

**Figure: Virtual Wire Deployment**

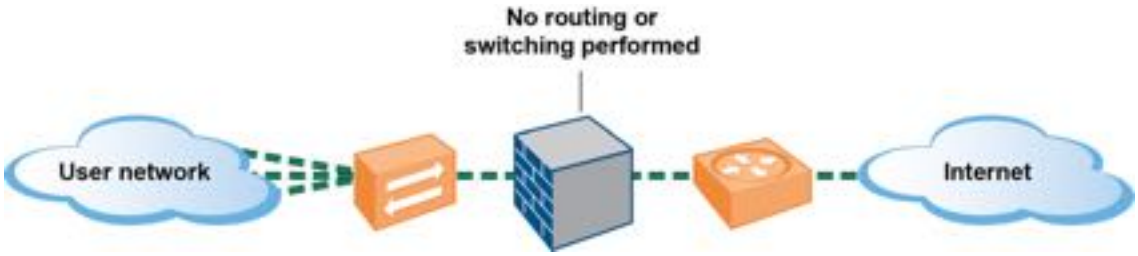


### Virtual Wire Subinterfaces

Virtual wire subinterfaces provide flexibility in enforcing distinct policies when you need to manage traffic from multiple customer networks. It allows you to separate and classify traffic into different zones (the zones can belong to separate virtual systems, if required) using the following criteria:

- **VLAN tags** —The example in Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only), shows an Internet Service Provider (ISP) using virtual wire subinterfaces with VLAN tags to separate traffic for two different customers.

 © Palo Alto Networks, Inc.

- **VLAN tags in conjunction with IP classifiers (address, range, or subnet)**— The following example shows an ISP with two separate virtual systems on a firewall that manages traffic from two different customers. On each virtual system, the example illustrates how virtual wire subinterfaces with VLAN tags and IP classifiers are used to classify traffic into separate zones and apply relevant policy for customers from each network.

| Virtual Wire Subinterface Workflow |
|---|
| 1. Configure two Ethernet interfaces as type virtual wire, and assign these interfaces to a virtual wire. |
| 2. Create subinterfaces on the parent Virtual Wire to separate CustomerA and CustomerB traffic. Make sure that the VLAN tags defined on each pair of subinterfaces that are configured as virtual wire(s) are identical. This is essential because a virtual wire does not switch VLAN tags. |
| 3. Create new subinterfaces and define IP classifiers. This task is optional and only required if you wish to add additional subinterfaces with IP classifiers for further managing traffic from a customer based on the combination of VLAN tags and a specific source IP address, range or subnet.<br>You can also use IP classifiers for managing untagged traffic. To do so, you must create a sub-interface with the vlan tag "0", and define sub-interface(s) with IP classifiers for managing untagged traffic using IP classifiers |

IP classification may only be used on the subinterfaces associated with one side of the virtual wire. The subinterfaces defined on the corresponding side of the virtual wire must use the same VLAN tag, but must not include an IP classifier.

**Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only)**

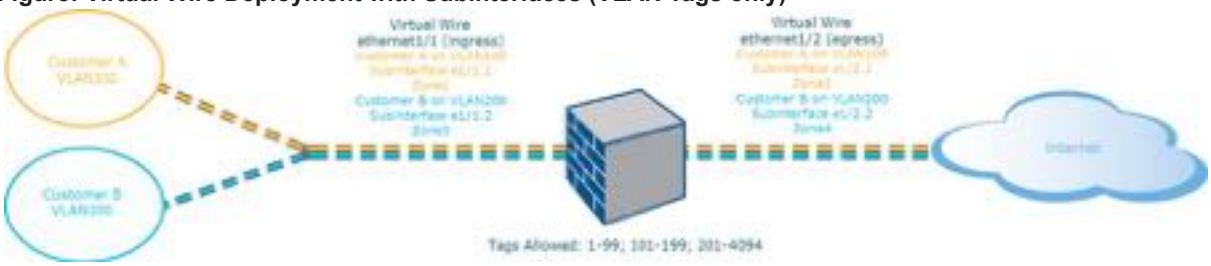


Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only) depicts CustomerA and CustomerB connected to the firewall through one physical interface, ethernet1/1, configured as a Virtual Wire; it is the ingress interface. A second physical interface, ethernet1/2, is also part of the Virtual Wire; it is the egress interface that provides access to the Internet. For CustomerA, you also have subinterfaces ethernet1/1.1 (ingress) and ethernet1/2.1 (egress). For CustomerB, you have the subinterface ethernet1/1.2 (ingress) and ethernet1/2.2 (egress). When configuring the subinterfaces, you must assign the appropriate VLAN tag and zone in order to apply policies for each customer. In this example, the policies for CustomerA are created between Zone1 and Zone2, and policies for CustomerB are created between Zone3 and Zone4.

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this example, a single subinterface matches the VLAN tag on the incoming packet, hence that subinterface is selected. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

The same VLAN tag must not be defined on the parent virtual wire interface and the subinterface. Verify that the VLAN tags defined on the Tag Allowed list of the parent virtual wire interface (**Network > Virtual Wires**) are not included on a subinterface.

Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers) depicts CustomerA and CustomerB connected to one physical firewall that has two virtual systems (vsys), in addition to the default virtual system (vsys1). Each virtual system is an independent virtual firewall that is managed separately for each customer. Each vsys has attached interfaces/subinterfaces and security zones that are managed independently.

**Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers)**

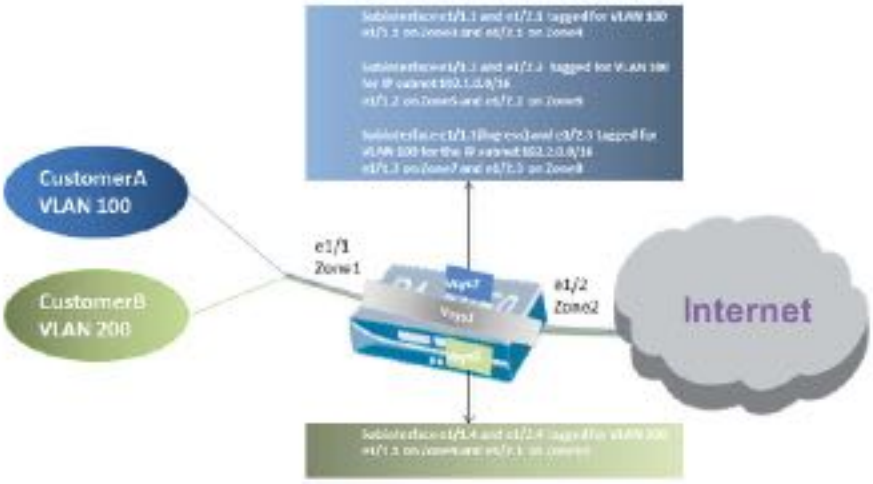


Vsys1 is set up to use the physical interfaces ethernet1/1 and ethernet1/2 as a virtual wire; ethernet1/1 is the ingress interface and ethernet1/2 is the egress interface that provides access to the Internet. This virtual wire is configured to accept all tagged and untagged traffic with the exception of VLAN tags 100 and 200 that are assigned to the subinterfaces.

CustomerA is managed on vsys2 and CustomerB is managed on vsys3. On vsys2 and vsys3, the following vwire subinterfaces are created with the appropriate VLAN tags and zones to enforce policy measures.

| Customer | Vsys | Vwire Subinterfaces | Zone | VLAN Tag | IP Classifier |
|---|---|---|---|---|---|
| A | 2 | e1/1.1 (ingress)<br>e1/2.1 (egress) | Zone3<br>Zone4 | 100<br>100 | None |
| | 2 | e1/1.2 (ingress)<br>e1/2.2 (egress) | Zone5<br>Zone6 | 100<br>100 | IP subnet<br>192.1.0.0/16 |
| | 2 | e1/1.3 (ingress)<br>e1/2.3 (egress) | Zone7<br>Zone8 | 100<br>100 | IP subnet<br>192.2.0.0/16 |
| B | 3 | e1/1.4 (ingress)<br>e1/2.4 (egress) | Zone9<br>Zone10 | 200<br>200 | None |

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this case, for CustomerA, there are multiple subinterfaces that use the same VLAN tag. Hence, the firewall first narrows the classification to a subinterface based on the source IP address in the packet. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

 © Palo Alto Networks, Inc.

For return-path traffic, the firewall compares the destination IP address as defined in the IP classifier on the customer-facing subinterface and selects the appropriate virtual wire to route traffic through the accurate subinterface.

The same VLAN tag must not be defined on the parent virtual wire interface and the subinterface. Verify that the VLAN tags defined on the Tag Allowed list of the parent virtual wire interface (**Network > Virtual Wires**) are not included on a subinterface.

## Layer 2 Deployments

In a Layer 2 deployment, the firewall provides switching between two or more networks. Each group of interfaces must be assigned to a VLAN object in order for the firewall to switch between them. The firewall will perform VLAN tag switching when layer 2 subinterfaces are attached to a common VLAN object. Choose this option when switching is required.

**Figure: Layer 2 Deployment**

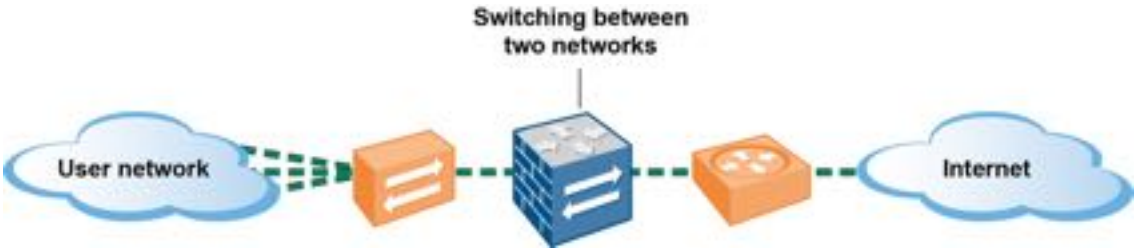


## Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between multiple ports. An IP address must be assigned to each interface and a virtual router must be defined to route the traffic. Choose this option when routing is required.

**Figure: Layer 3 Deployment**

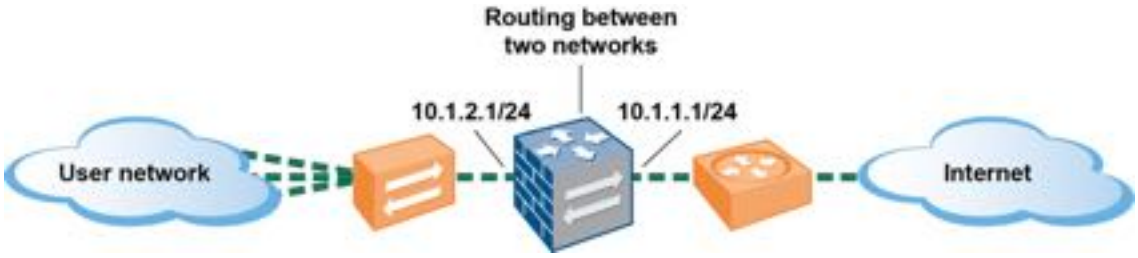


In addition, because the firewall must route traffic in a Layer 3 deployment, you must configure a virtual router. See Virtual Routers.

### Point-to-Point Protocol over Ethernet Support

You can configure the firewall to be a Point-to-Point Protocol over Ethernet (PPPoE) termination point to support connectivity in a Digital Subscriber Line (DSL) environment where there is a DSL modem but no other PPPoE device to terminate the connection.

You can choose the PPPoE option and configure the associated settings when an interface is defined as a Layer 3 interface.

PPPoE is not supported in HA active/active mode.

### DHCP Client

You can configure the firewall interface to act as a DHCP client and receive a dynamically assigned IP address. The firewall also provides the capability to propagate settings received by the DHCP client interface into a DHCP server operating on the firewall. This is most commonly used to propagate DNS server settings from an Internet service provider to client machines operating on the network protected by the firewall.

DHCP client is not supported in HA active/active mode.

For more information, see DHCP.

## Tap Mode Deployments

A network tap is a device that provides a way to access data flowing across a computer network. Tap mode deployment allows you to passively monitor traffic flows across a network by way of a switch SPAN or mirror port.

The SPAN or mirror port permits the copying of traffic from other ports on the switch. By dedicating an interface on the firewall as a tap mode interface and connecting it with a switch SPAN port, the switch SPAN port provides the firewall with the mirrored traffic. This provides application visibility within the network without being in the flow of network traffic.

When deployed in tap mode, the firewall is not able to take action, such as block traffic or apply QoS traffic control.

 © Palo Alto Networks, Inc.

# Virtual Routers

The firewall uses virtual routers to obtain routes to other subnets by manually defining a route (static routes) or through participation in Layer 3 routing protocols (dynamic routes). The best routes obtained through these methods are used to populate the firewall's IP route table. When a packet is destined for a different subnet, the Virtual Router obtains the best route from this IP route table and forwards the packet to the next hop router defined in the table.

The Ethernet interfaces and VLAN interfaces defined on the firewall receive and forward the Layer 3 traffic. The destination zone is derived from the outgoing interface based on the forwarding criteria, and policy rules are consulted to identify the security policies to be applied. In addition to routing to other network devices, virtual routers can route to other virtual routers within the same firewall if a next hop is specified to point to another virtual router.

You can configure the virtual router to participate with dynamic routing protocols (BGP, OSPF, or RIP) as well as adding static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that are not shared between virtual routers, enabling you to configure different routing behaviors for different interfaces.

Each Layer 3 interface, loopback interface, and VLAN interface defined on the firewall must be associated with a virtual router. While each interface can belong to only one virtual router, multiple routing protocols and static routes can be configured for a virtual router. Regardless of the static routes and dynamic routing protocols configured for a virtual router, a common general configuration is required. The firewall uses Ethernet switching to reach other devices on the same IP subnet.

The following Layer 3 routing protocols are supported from Virtual Routers:

- RIP
- OSPF
- OSPFv3
- BGP

| **Define a Virtual Router General Configuration** | | |
|---|---|---|
| Step 1 | Gather the required information from your network administrator. | • Interfaces that you want to route<br>• Administrative distances for static, OSFP internal, OSPF external, IBGP, EBGP and RIP |
| Step 2 | Create the virtual router and name it. | 1. Select **Network > Virtual Routers**.<br>2. Click **Add** and enter a name for the virtual router.<br>3. Select interfaces to apply to the virtual router.<br>4. Click **OK**. |
| Step 3 | Select interfaces to apply to the virtual router. | 1. Click **Add** in the Interfaces box.<br>2. Select an already defined interface from the drop-down.<br>3. Repeat Step 2 for all interfaces that you want to add to the virtual router. |

© Palo Alto Networks, Inc.

| Define a Virtual Router General Configuration (Continued) | | |
|---|---|---|
| Step 4 | Set Administrative Distances for static and dynamic routing. | 1. Set Administrative Distances as required.<br>• **Static**—Range: 10-240, Default: 10<br>• **OSPF Internal**—Range: 10-240, Default: 30<br>• **OSPF External**—Range: 10-240, Default: 110<br>• **IBGP**—Range: 10-240, Default: 200<br>• **EBGP**—Range: 10-240, Default: 20<br>• **RIP**—Range: 10-240, Default: 120 |
| Step 5 | Save virtual router general settings. | Click **OK** to save your settings. |
| Step 6 | Commit your changes. | Click **Commit**. The device may take up to 90 seconds to save your changes. |

 © Palo Alto Networks, Inc.

# Static Routes

The following procedure shows how to integrate the firewall into the network using static routing.

| Set Up Interfaces and Zones | | |
|---|---|---|
| Step 1 | Configure a default route to your Internet router. | 1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog.<br>2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and enter the route in the **Destination** field (for example, 0.0.0.0/0).<br>3. Select the **IP Address** radio button in the **Next Hop** field and then enter the IP address and netmask for your Internet gateway (for example, 208.80.56.1).<br>4. Click **OK** twice to save the virtual router configuration. |
| Step 2 | Configure the external interface (the interface that connects to the Internet). | 1. Select **Network > Interfaces** and then select the interface you want to configure. In this example, we are configuring Ethernet1/3 as the external interface.<br>2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.<br>3. In the **Virtual Router** drop-down, select **default**.<br>4. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Untrust, and then click **OK**.<br>5. To assign an IP address to the interface, select the **IPv4** tab and **Static** radio button. Click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 208.80.56.100/24.<br>6. To enable you to ping the interface, select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. Enter a **Name** for the profile, select **Ping** and then click **OK**.<br>7. To save the interface configuration, click **OK**. |

| Set Up Interfaces and Zones (Continued) | | |
|---|---|---|
| Step 3 | Configure the interface that connects to your internal network.<br><br>In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT. See Configure NAT for details. | 1. Select **Network > Interfaces** and select the interface you want to configure. In this example, we are configuring Ethernet1/4 as the internal interface.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Trust, and then click **OK**.<br>4. Select the same Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab and the **Static** radio button, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 4 | Configure the interface that connects to the DMZ. | 1. Select the interface you want to configure.<br>2. Select **Layer3** from the **Interface Type** drop-down. In this example, we are configuring Ethernet1/13 as the DMZ interface.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example DMZ, and then click **OK**.<br>4. Select the Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab and the **Static** radio button, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 5 | Save the interface configuration. | Click **Commit**. |
| Step 6 | Cable the firewall. | Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment. |
| Step 7 | Verify that the interfaces are active. | From the web interface, select **Network > Interfaces** and verify that icon in the Link State column is green. You can also monitor link state from the **Interfaces** widget on the **Dashboard**.<br><br> 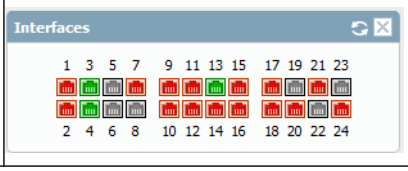 |

 © Palo Alto Networks, Inc.

# RIP

Routing Information Protocol (RIP) is an interior gateway protocol (IGP) that was designed for small IP networks. RIP relies on hop count to determine routes; the best routes have the fewest number of hops. RIP is based on UDP and uses port 520 for route updates. By limiting routes to a maximum of 15 hops, the protocol helps prevent the development of routing loops, but also limits the supported network size. If more than 15 hops are required, traffic is not routed. RIP also can take longer to converge than OSPF and other routing protocols. The firewall supports RIP v2.

Perform the following procedure to configure RIP.

| **Configure RIP** | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Virtual Routers for details. |
| Step 2 | Configure general RIP configuration settings. | 1. Select the **RIP** tab.<br>2. Select the **Enable** check box to enable the RIP protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through RIP. This is the recommended default setting.<br>4. De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through RIP. |
| Step 3 | Configure interfaces for the RIP protocol. | 1. On **Interfaces** tab, select an interface from the drop-down in the Interface configuration section.<br>2. Select an already defined interface from the drop-down.<br>3. Select the **Enable** check box.<br>4. Select the **Advertise** check box to advertise a default route to RIP peers with the specified metric value.<br>5. You can optionally select a profile from the **Auth Profile** drop-down. See Step 5 for details.<br>6. Select normal, passive or send-only from the **Mode** drop-down.<br>7. Click **OK**. |
| Step 4 | Configure RIP timers. | 1. On the **Timers** tab, enter a value in the **Interval Seconds (sec)** box. This setting defines the length of the timer interval in seconds. This duration is used for the remaining RIP timing fields. Default: 1. Range: 1 - 60.<br>2. Enter an **Update Intervals** to define the number of intervals between route update announcements Default: 30. Range: 1 - 3600.<br>3. Specify the **Delete Intervals** to define the number of intervals between the time that the route expires to its deletion Default: 180. Range: 1 - 3600.<br>4. Specify the **Expire Intervals** to define the number of intervals between the time that the route was last updated to its expiration Default: 120. Range: 1 - 3600. |

| Configure RIP | |
|---|---|
| Step 5 (Optional) Configure Auth Profiles. | By default, the firewall does not use RIP authentication for the exchange between RIP neighbors. Optionally, you can configure RIP authentication between RIP neighbors by either a simple password or using MD5 authentication.<br>**Simple Password RIP authentication**<br>1. Select **Auth Profiles** and click **Add**.<br>2. Enter a name for the authentication profile to authenticate RIP messages.<br>3. Select **Simple Password** as the **Password Type**.<br>4. Enter a simple password and then confirm.<br>**MD5 RIP authentication**<br>1. Select **Auth Profiles** and click **Add**.<br>2. Enter a name for the authentication profile to authenticate RIP messages.<br>3. Select **MD5** as the **Password Type**.<br>4. Click **Add**.<br>5. Enter one or more password entries, including:<br>• Key-ID Range 0-255<br>• Key<br>6. You can optionally select **Preferred** status.<br>7. Click **OK** to specify the key to be used to authenticate outgoing message.<br>8. Click **OK** again in the Virtual Router - RIP Auth Profile dialog box. |

 © Palo Alto Networks, Inc.

# OSPF

Open Shortest Path First (OSPF) is an interior gateway protocol (IGP) that is most often used to dynamically manage network routes in large enterprise network. It determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The information gathered from the LSAs is used to construct a topology map of the network. This topology map is shared across routers in the network and used to populate the IP routing table with available routes.

Changes in the network topology are detected dynamically and used to generate a new topology map within seconds. A shortest path tree is computed of each route. Metrics associated with each routing interface are used to calculate the best route. These can include distance, network throughput, link availability etc. Additionally, these metrics can be configured statically to direct the outcome of the OSPF topology map.

Palo Alto networks implementation of OSPF fully supports the following RFCs:

- RFC 2328 (for IPv4)
- RFC 5340 (for IPv6)

The following topics provide more information about the OSPF and procedures for configuring OSPF on the firewall:

▲ OSPF Concepts
▲ Configure OSPF
▲ Configure OSPFv3
▲ Configure OSPF Graceful Restart
▲ Confirm OSPF Operation

## OSPF Concepts

The following topics introduce the OSPF concepts you will need to understand in order to configure the firewall to participate in an OSPF network:

▲ OSPFv3
▲ OSPF Neighbors
▲ OSPF Areas
▲ OSPF Router Types

© Palo Alto Networks, Inc.

## OSPFv3

OSPFv3 provides support for the OSPF routing protocol within an IPv6 network. As such, it provides support for IPv6 addresses and prefixes. It retains most of the structure and functions in OSPFv2 (for IPv4) with some minor changes. The following are some of the additions and changes to OSPFv3:

- **Support for multiple instances per link**—With OSPFv3, you can run multiple instances of the OSPF protocol over a single link. This is accomplished by assigning an OSPFv3 instance ID number. An interface that is assigned to an instance ID drops packets that contain a different ID.
- **Protocol Processing Per-link**—OSPFv3 operates per-link instead of per-IP-subnet as on OSPFv2.
- **Changes to Addressing**—IPv6 addresses are not present in OSPFv3 packets, except for LSA payloads within link state update packets. Neighboring routers are identified by the Router ID.
- **Authentication Changes**—OSPFv3 doesn't include any authentication capabilities. Configuring OSPFv3 on a firewall requires an authentication profile that specifies Encapsulating Security Payload (ESP) or IPv6 Authentication Header (AH).The re-keying procedure specified in RFC 4552 is not supported in this release.
- **Support for multiple instances per-link**—Each instance corresponds to an instance ID contained in the OSPFv3 packet header.
- **New LSA Types**—OSPFv3 supports two new LSA types: Link LSA and Intra Area Prefix LSA.

All additional changes are described in detail in RFC 5340.

## OSPF Neighbors

Two OSPF-enabled routers connected by a common network and in the same OSPF area that form a relationship are OSPF neighbors. The connection between these routers can be through a common broadcast domain or by a point-to-point connection. This connection is made through the exchange of hello OSPF protocol packets. These neighbor relationships are used to exchange routing updates between routers.

 © Palo Alto Networks, Inc.

## OSPF Areas

OSPF operates within a single autonomous system (AS). Networks within this single AS, however, can be divided into a number of areas. By default, Area 0 is created. Area 0 can either function alone or act as the OSPF backbone for a larger number of areas. Each OSPF area is named using a 32-bit identifier which in most cases is written in the same dotted-decimal notation as an IP4 address. For example, Area 0 is usually written as 0.0.0.0.

The topology of an area is maintained in its own link state database and is hidden from other areas, which reduces the amount of traffic routing required by OSPF. The topology is then shared in a summarized form between areas by a connecting router.

**Table: OSPF Area Types**

| Area Type | Description |
|---|---|
| Backbone Area | The backbone area (Area 0) is the core of an OSPF network. All other areas are connected to it and all traffic between areas must traverse it. All routing between areas is distributed through the backbone area. While all other OSPF areas must connect to the backbone area, this connection doesn't need to be direct and can be made through a virtual link. |
| Normal OSPF Area | In a normal OSPF area there are no restrictions; the area can carry all types of routes. |
| Stub OSPF Area | A stub area does not receive routes from other autonomous systems. Routing from the stub area is performed through the default route to the backbone area. |
| NSSA Area | The Not So Stubby Area (NSSA) is a type of stub area that can import external routes, with some limited exceptions. |

## OSPF Router Types

Within an OSPF area, routers are divided into the following categories.

- **Internal Router**—A router with that has OSPF neighbor relationships only with devices in the same area.
- **Area Border Router (ABR)**—A router that has OSPF neighbor relationships with devices in multiple areas. ABRs gather topology information from their attached areas and distribute it to the backbone area.
- **Backbone Router**—A backbone router is any OSPF router that is attached to the OSPF backbone. Since ABRs are always connected to the backbone, they are always classified as backbone routers.
- **Autonomous System Boundary Router (ASBR)**—An ASBR is a router that attaches to more than one routing protocol and exchanges routing information between them.

© Palo Alto Networks, Inc.

## Configure OSPF

OSPF determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The router keeps information about the links between it and the destination and can make highly efficient routing decisions. A cost is assigned to each router interface, and the best routes are determined to be those with the lowest costs, when summed over all the encountered outbound router interfaces and the interface receiving the LSA.

Hierarchical techniques are used to limit the number of routes that must be advertised and the associated LSAs. Because OSPF dynamically processes a considerable amount of route information, it has greater processor and memory requirements than does RIP.

| **Configure OSPF** | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Virtual Routers for details. |
| Step 2 | Configure general OSPF configuration settings. | 1. Select the **OSPF** tab.<br>2. Select the **Enable** check box to enable the OSPF protocol.<br>3. Optionally enter the **Router ID**.<br>4. Select the **Reject Default Route** check box if you do not want to learn any default routes through OSPF. This is the recommended default setting.<br>De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through OSPF. |


© Palo Alto Networks, Inc.

**Configure OSPF (Continued)**

| | |
|---|---|
| Step 3 Configure Areas - Type for the OSFP protocol. | 1. On the **Areas** tab, click **Add**.<br>2. Enter an Area ID for the area in *x.x.x.x* format. This is the identifier that each neighbor must accept to be part of the same area.<br>3. On the **Type** tab, select one of the following from the area **Type** drop-down:<br>• **Normal**—There are no restrictions; the area can carry all types of routes.<br>• **Stub**—There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, configure the following:<br>– **Accept Summary**—Link state advertisements (LSA) are accepted from other areas. If this option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.<br>– **Advertise Default Route**—Default route LSAs will be included in advertisements to the stub area along with a configured metric value in the configured range: 1-255.<br>• **NSSA** (Not-So-Stubby Area)—The firewall can only leave the area by routes other than OSPF routes. If selected, configure **Accept Summary** and **Advertise Default Route** as described for **Stub**. If you select this option, configure the following:<br>– **Type**—Select either **Ext 1** or **Ext 2** route type to advertise the default LSA.<br>– **Ext Ranges**—Click **Add** in the section to enter ranges of external routes that you want to enable or suppress advertising for.<br>4. **Priority**—Enter the OSPF priority for this interface (0-255). This is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR) according to the OSPF protocol. When the value is zero, the router will not be elected as a DR or BDR.<br>• **Auth Profile**—Select a previously-defined authentication profile.<br>• **Timing**—It is recommended that you keep the default timing settings.<br>• **Neighbors**—For p2pmp interfaces, enter the neighbor IP address for all neighbors that are reachable through this interface.<br>5. Select **normal**, **passive** or **send-only** as the **Mode**.<br>6. Click **OK**. |

| Configure OSPF (Continued) | | |
|---|---|---|
| Step 4 | Configure Areas - Range for the OSFP protocol | 1. On the **Range** tab, click **Add** to aggregate LSA destination addresses in the area into subnets.<br>2. **Advertise** or **Suppress** advertising LSAs that match the subnet, and click **OK**. Repeat to add additional ranges. |
| Step 5 | Configure Areas - Interfaces for the OSFP protocol | 1. On the **Interface** tab, click **Add** and enter the following information for each interface to be included in the area:<br>• **Interface**—Select an interface from the drop-down box.<br>• **Enable**—Selecting this option causes the OSPF interface settings to take effect.<br>• **Passive**—Select the check box to if you do not want the OSPF interface to send or receive OSPF packets. Although OSPF packets are not sent or received if you choose this option, the interface is included in the LSA database.<br>• **Link type**—Choose **Broadcast** if you want all neighbors that are accessible through the interface to be discovered automatically by multicasting OSPF hello messages, such as an Ethernet interface. Choose **p2p** (point-to-point) to automatically discover the neighbor. Choose **p2mp** (point-to-multipoint) when neighbors must be defined manually. Defining neighbors manually is allowed only for **p2mp** mode.<br>• **Metric**—Enter an OSPF metric for this interface. Default: 10. Range: 0-65535.<br>• **Priority**—Enter an OSPF priority for this interface. This is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR). Default: 1. Range: 0 - 255. If zero is configured, the router will not be elected as a DR or BDR.<br>• **Auth Profile**—Select a previously-defined authentication profile.<br>• **Timing**—The following OSPF timing settings can be set. Palo Alto Networks recommends that you retain the default timing settings.<br>– **Hello Interval (sec)**—Interval at which the OSPF process sends hello packets to its directly connected neighbors. Default: 10 seconds. Range: 0-3600 seconds.<br>– **Dead Counts—**Number of times the hello interval can occur for a neighbor without OSPF receiving a hello packet from the neighbor, before OSPF considers that neighbor down. The **Hello Interval** multiplied by the **Dead Counts** equals the value of the dead timer. Default: 4. Range: 3-20. |


© Palo Alto Networks, Inc.

| Configure OSPF (Continued) | |
|---|---|
| | – **Retransmit Interval (sec)**—Length of time that OSPF waits to receive a link state advertisement (LSA) from a neighbor before OSPF retransmits the LSA. Default: 10 seconds. Range: 0-3600 seconds.<br>– **Transit Delay (sec)**—Length of time that an LSA is delayed before it is sent out of an interface. Default: 1 second. Range: 0-3600 seconds.<br>– **Graceful Restart Hello Delay (sec)**—Applies to an OSPF interface when Active/Passive High Availability is configured. **Graceful Restart Hello Delay** is the length of time during which the firewall sends Grace LSA packets at 1-second intervals. During this time, no hello packets are sent from the restarting firewall. During the restart, the dead timer (which is the **Hello Interval** multiplied by the **Dead Counts**) is also counting down. If the dead timer is too short, the adjacency will go down during the graceful restart because of the hello delay. Therefore, it is recommended that the dead timer be at least four times the value of the **Graceful Restart Hello Delay**. For example, a **Hello Interval** of 10 seconds and a **Dead Counts** of 4 yield a dead timer of 40 seconds. If the **Graceful Restart Hello Delay** is set to 10 seconds, that 10-second delay of hello packets is comfortably within the 40-second dead timer, so the adjacency will not time out during a graceful restart. Default: 10 seconds. Range: 1-10 seconds.<br>• If **p2mp** is selected for **Link Type** interfaces, enter the neighbor IP addresses for all neighbors that are reachable through this interface.<br>2. Click **OK** |
| Step 6 Configure Areas - Virtual Links. | 1. On the **Virtual Link** tab, click **Add** and enter the following information for each virtual link to be included in the backbone area:<br>• **Name**—Enter a name for the virtual link.<br>• **Neighbor ID**—Enter the router ID of the router (neighbor) on the other side of the virtual link.<br>• **Transit Area**—Enter the area ID of the transit area that physically contains the virtual link.<br>• **Enable**—Select to enable the virtual link.<br>• **Timing**—It is recommended that you keep the default timing settings.<br>• **Auth Profile**—Select a previously-defined authentication profile.<br>2. Click **OK**. |

| Configure OSPF (Continued) | |
|---|---|
| Step 7 (Optional) Configure Auth Profiles. | By default, the firewall does not use OSPF authentication for the exchange between OSPF neighbors. Optionally, you can configure OSPF authentication between OSPF neighbors by either a simple password or using MD5 authentication.<br>**Simple Password OSPF authentication**<br>1. On the **Auth Profiles** tab, click **Add**.<br>2. Enter a name for the authentication profile to authenticate OSPF messages.<br>3. Select **Simple Password** as the **Password Type**.<br>4. Enter a simple password and then confirm.<br>**MD5 OSPF authentication**<br>1. On the **Auth Profiles** tab, click **Add**.<br>2. Enter a name for the authentication profile to authenticate OSPF messages.<br>3. Select **MD5** as the **Password Type**.<br>4. Click **Add**.<br>5. Enter one or more password entries, including:<br>• Key-ID Range 0-255<br>• Key<br>• Select the **Preferred** option to specify that the key be used to authenticate outgoing messages.<br>6. Click **OK**.<br>7. Click **OK** again in the Virtual Router - OSPF Auth Profile dialog box. |
| Step 8 Configure Advanced OSPF options. | 1. On the **Advanced** tab, select the **RFC 1583 Compatibility** check box to ensure compatibility with RFC 1583.<br>2. Configure a value for the **SPF Calculation Delay** (sec) timer.<br>This timer allows you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times.<br>3. Configure a value for the **LSA Interval (sec)** time. This timer specifies the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur. |

 © Palo Alto Networks, Inc.

# Configure OSPFv3

| Configure OSPFv3 | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Virtual Routers for details. |
| Step 2 | Configure general OSPF configuration settings. | 1. Select the **OSPF** tab.<br>2. Select the **Enable** check box to enable the OSPF protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through OSPF. This is the recommended default setting.<br>4. Clear the **Reject Default Route** check box if you want to permit redistribution of default routes through OSPF. |
| Step 3 | Configure general OSPFv3 configuration settings. | 1. Select the **OSPFv3** tab.<br>2. Select the **Enable** check box to enable the OSPF protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through OSPFv3 This is the recommended default setting.<br>De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through OSPFv3. |
| Step 4 | Configure Auth Profile for the OSFPv3 protocol.<br>While OSPFv3 doesn't include any authentication capabilities of its own, it relies entirely on IPsec to secure communications between neighbors. | When configuring an authentication profile, you must use Encapsulating Security Payload (ESP) or IPv6 Authentication Header (AH).<br>**ESP OSPFv3 authentication**<br>1. On the **Auth Profiles** tab, click **Add**.<br>2. Enter a name for the authentication profile to authenticate OSPFv3 messages.<br>3. Specify a Security Policy Index (**SPI**). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a hexadecimal value between 00000000 and FFFFFFFF.<br>4. Select **ESP** for **Protocol**.<br>5. Select a **Crypto Algorithm** from the drop-down.<br>You can enter none or one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5.<br>6. If a **Crypto Algorithm** other than none was selected, enter a value for **Key** and then confirm. |

| Configure OSPFv3 (Continued) | |
|---|---|
| | **AH OSPFv3 authentication**<br>1. On the **Auth Profiles** tab, click **Add**.<br>2. Enter a name for the authentication profile to authenticate OSPFv3 messages.<br>3. Specify a Security Policy Index (**SPI**). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a hexadecimal value between 00000000 and FFFFFFFF.<br>4. Select **AH** for **Protocol**.<br>5. Select a **Crypto Algorithm** from the drop-down.<br>You must enter one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5.<br>6. Enter a value for **Key** and then confirm.<br>7. Click **OK**.<br>8. Click **OK** again in the Virtual Router - OSPF Auth Profile dialog. |
| Step 5 Configure Areas - Type for the OSFP protocol. | 1. On the **Areas** tab, click **Add**.<br>2. Enter an Area ID. This is the identifier that each neighbor must accept to be part of the same area.<br>3. On the **General** tab, select one of the following from the area **Type** drop-down:<br>• **Normal**—There are no restrictions; the area can carry all types of routes.<br>• **Stub**—There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, configure the following:<br>– **Accept Summary**—Link state advertisements (LSA) are accepted from other areas. If this option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.<br>– **Advertise Default Route**—Default route LSAs will be included in advertisements to the stub area along with a configured metric value in the configured range: 1-255. |

 © Palo Alto Networks, Inc.

| Configure OSPFv3 (Continued) | | |
|---|---|---|
| | | • **NSSA** (Not-So-Stubby Area)—The firewall can only leave the area by routes other than OSPF routes. If selected, configure **Accept Summary** and **Advertise Default Route** as described for **Stub**. If you select this option, configure the following:<br>– **Type**—Select either **Ext 1** or **Ext 2** route type to advertise the default LSA.<br>– **Ext Ranges**—Click **Add** in the section to enter ranges of external routes that you want to enable or suppress advertising for. |
| Step 6 | Associate an OSPFv3 authentication profile to an area or an interface. | **To an Area**<br>1. On the **Areas** tab, select an existing area from the table.<br>2. On the **General** tab, select a previously defined **Authentication Profile** from the **Authentication** drop-down.<br>3. Click **OK**.<br>**To an Interface**<br>1. On the **Areas** tab, select an existing area from the table.<br>2. Select the **Interface** tab and click **Add**.<br>3. Select the authentication profile you want to associate with the OSPF interface from the **Auth Profile** drop-down. |
| Step 7 | (Optional) Configure Export Rules | 1. On the **Export** tab, click **Add**.<br>2. Select the **Allow Redistribute Default Route** check box to permit redistribution of default routes through OSPFv3.<br>3. Select the name of a redistribution profile. The value must be an IP subnet or valid redistribution profile name.<br>4. Select a metric to apply for **New Path Type**.<br>5. Specify a **New Tag** for the matched route that has a 32-bit value.<br>6. Assign a metric for the new rule. Range: 1 - 65535.<br>7. Click **OK**. |

| Configure OSPFv3 (Continued) | |
|---|---|
| Step 8 Configure Advanced OSPFv3 options. | 1. On the **Advanced** tab, select the **Disable Transit Routing for SPF Calculation** check box if you want the firewall to participate in OSPF topology distribution without being used to forward transit traffic.<br>2. Configure a value for the **SPF Calculation Delay** (sec) timer.<br>This timer allows you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times.<br>3. Configure a value for the **LSA Interval (sec) time**, This timer specifies the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur.<br>4. Optionally, Configure OSPF Graceful Restart. |

## Configure OSPF Graceful Restart

OSPF Graceful Restart directs OSPF neighbors to continue using routes through a device during a short transition when it is out of service. This behavior increases network stability by reducing the frequency of routing table reconfiguration and the related route flapping that can occur during short periodic down times.

For a Palo Alto Networks firewall, OSPF Graceful Restart involves the following operations:

- **Firewall as a restarting device**—In a situation where the firewall will be down for a short period of time or is unavailable for short intervals, it sends Grace LSAs to its OSPF neighbors. The neighbors must be configured to run in Graceful Restart Helper mode. In Helper Mode, the neighbors receive the Grace LSAs that inform it that the firewall will perform a graceful restart within a specified period of time defined as the **Grace Period**. During the grace period, the neighbor continues to forward routes through the firewall and to send LSAs that announce routes through the firewall. If the firewall resumes operation before expiration of the grace period, traffic forwarding will continue as before without network disruption. If the firewall does not resume operation after the grace period has expired, the neighbors will exit helper mode and resume normal operation, which will involve reconfiguring the routing table to bypass the firewall.
- **Firewall as a Graceful Restart Helper**—In a situation where neighboring routers may be down for a short periods of time, the firewall can be configured to operate in Graceful Restart Helper mode. If configured in this mode, the firewall will be configured with a **Max Neighbor Restart Time**. When the firewall receives the Grace LSAs from its OSFP neighbor, it will continue to route traffic to the neighbor and advertise routes through the neighbor until either the grace period or max neighbor restart time expires. If neither expires before the neighbor returns to service, traffic forwarding continues as before without network disruption. If either period expires before the neighbor returns to service, the firewall will exit helper mode and resume normal operation, which will involve reconfiguring the routing table to bypass the neighbor.

 © Palo Alto Networks, Inc.

| Configure OSPF Graceful Restart |
|---|
| 1. Select **Network > Virtual Routers** and select the virtual router you want to configure. |
| 2. Select **OSPF > Advanced**. |
| 3. Verify that the following check boxes are selected (they are enabled by default):<br>• **Enable Graceful Restart**<br>• **Enable Helper Mode**<br>• **Enable Strict LSA checking**<br>These check boxes should remain selected unless required by your topology. |
| 4. Configure a **Grace Period** in seconds. |
| 5. Configure a **Max Neighbor Restart Time** in seconds. |

## Confirm OSPF Operation

Once an OSPF configuration has been committed, you can use any of the following operations to confirm that OSPF is operating:

- View the Routing Table
- Confirm OSPF Adjacencies
- Confirm that OSPF Connections are Established

### View the Routing Table

By viewing the routing table, you can see whether OSPF routes have been established. The routing table is accessible from either the web interface or the CLI. If you are using the CLI, use the following commands:

- `show routing route`
- `show routing fib`

The following procedure describes how to use the web interface to view the routing table.

| View the Routing Table |
|---|
| 1. Select **Network > Virtual Routers**. |
| 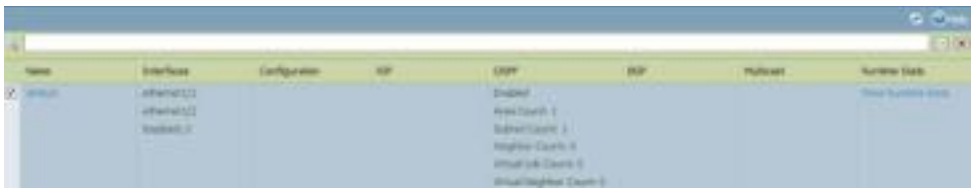 |

**View the Routing Table**

2. Select the **Routing** tab and examine the **Flags** column of the routing table for routes that were learned by OSPF.

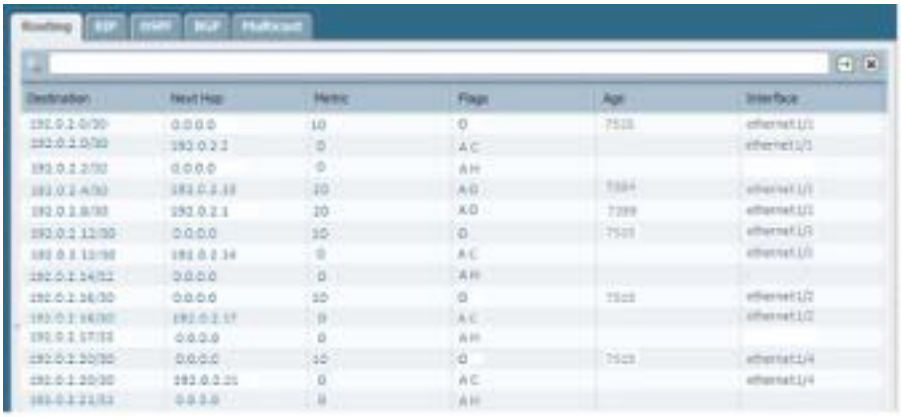

## Confirm OSPF Adjacencies

By viewing the **Neighbor** tab as described in the following procedure, you can confirm that OSPF adjacencies have been established.

**View the Neighbor Tab to Confirm OSPF Adjacencies**

1. Select **Network > Virtual Routers**.

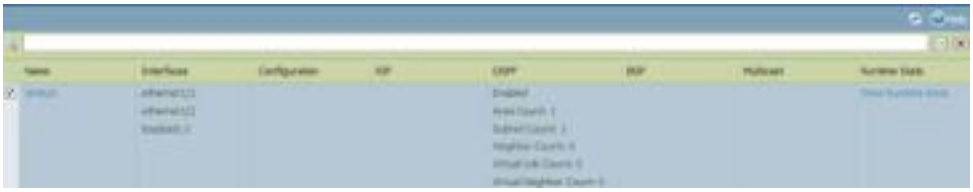

2. Select **OSPF > Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established.

 © Palo Alto Networks, Inc.

## Confirm that OSPF Connections are Established

By viewing the system log, you can confirm that OSPF connections have been established, as described in the following procedure:

**Examine the System Log**

1. Select **Monitor> System** and look for messages to confirm that OSPF adjacencies have been established.

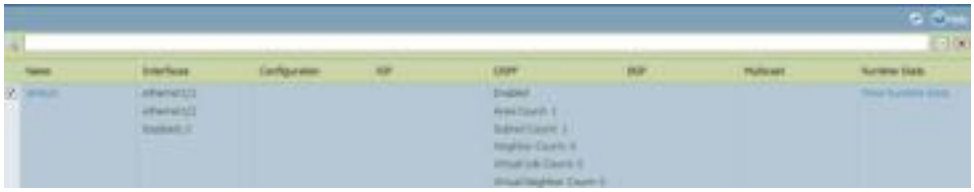

2. Select **OSPF > Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established.

# BGP

Border Gateway Protocol (BGP) is the primary Internet routing protocol. BGP determines network reachability based on IP prefixes that are available within autonomous systems (AS), where an AS is a set of IP prefixes that a network provider has designated to be part of a single routing policy.

In the routing process, connections are established between BGP peers (or neighbors). If a route is permitted by the policy, it is stored in the routing information base (RIB). Each time the local firewall RIB is updated, the firewall determines the optimal routes and sends an update to the external RIB, if export is enabled.

Conditional advertisement is used to control how BGP routes are advertised. The BGP routes must satisfy conditional advertisement rules before being advertised to peers.

BGP supports the specification of aggregates, which combine multiple routes into a single route. During the aggregation process, the first step is to find the corresponding aggregation rule by performing a longest match that compares the incoming route with the prefix values for other aggregation rules.

The firewall provides a complete BGP implementation, which includes the following features:

- Specification of one BGP routing instance per virtual router.
- Routing policies based on route-map to control import, export and advertisement, prefix-based filtering, and address aggregation.
- Advanced BGP features that include route reflector, AS confederation, route flap dampening, and graceful restart.
- IGP-BGP interaction to inject routes to BGP using redistribution profiles.

BGP configuration consists of the following elements:

- Per-routing-instance settings, which include basic parameters such as local route ID and local AS and advanced options such as path selection, route reflector, AS confederation, route flap, and dampening profiles.
- Authentication profiles, which specify the MD5 authentication key for BGP connections.
- Peer group and neighbor settings, which include neighbor address and remote AS and advanced options such as neighbor attributes and connections.
- Routing policy, which specifies rule sets that peer groups and peers use to implement imports, exports, conditional advertisements, and address aggregation controls.

Perform the following procedure to configure BGP.

| **Configure BGP** | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Virtual Routers for details. |

 © Palo Alto Networks, Inc.

| Configure BGP (Continued) | | |
|---|---|---|
| Step 2 | Configure standard BGP configuration settings. | 1. Select the **BGP** tab.<br>2. Select the **Enable** check box to enable the BGP protocol.<br>3. Assign an IP address to the virtual router in the **Router ID** box.<br>4. Enter the number of the AS to which the virtual router belongs in the **AS Number** box, based on the router ID. Range: 1-4294967295 |
| Step 3 | Configure general BGP configuration settings. | 1. Select **BGP > General**.<br>2. Select the **Reject Default Route** check box to ignore any default routes that are advertised by BGP peers.<br>3. Select the **Install Route** check box to install BGP routes in the global routing table.<br>4. Select the **Aggregate MED** check box to enable route aggregation even when routes have different Multi-Exit Discriminator (MED) values.<br>5. Enter a value for the **Default Local Preference** that specifies a value than can be used to determine preferences among different paths.<br>6. Select one of the following values for the AS format for interoperability purposes:<br>• 2 Byte (default value)<br>• 4 Byte<br>7. Enable or disable each of the following values for **Path Selection**:<br>• **Always Compare MED**—Enable this comparison to choose paths from neighbors in different autonomous systems.<br>• **Deterministic MED Comparison**—Enable this comparison to choose between routes that are advertised by IBGP peers (BGP peers in the same autonomous system).<br>8. Click **Add** to include a new authentication profile and configure the following settings:<br>• **Profile Name**—Enter a name to identify the profile.<br>• **Secret/Confirm Secret**—Enter and confirm a passphrase for BGP peer communications. |

| Configure BGP (Continued) | |
|---|---|
| Step 4 Configure BGP Advanced settings (Optional) | 1. On the **Advanced** tab, select **Graceful Restart** and configure the following timers:<br>• **Stale Route Time (sec)**—Specifies the length of time in seconds that a route can say in the stale state. Range: 1 - 3600 seconds. Default: 120 seconds.<br>• **Local Restart Time (sec)**—Specifies the length of time in seconds that the local device waits to restart. This value is advertised to peers. Range: 1 - 3600 seconds. Default: 120 seconds.<br>• **Max Peer Restart Time (sec)**—Specifies the maximum length of time in seconds that the local device accepts as a grave period restart time for peer devices. Range: 1 - 3600 seconds. Default: 120 seconds.<br>2. Specify an IPv4 identifier to represent the reflector cluster in the **Reflector Cluster ID** box.<br>3. Specify the identifier for the AS confederation to be presented as a single AS to external BGP peers in the **Confederation Member AS** box.<br>4. Click **Add** and enter the following information for each Dampening Profile that you want to configure, select **Enable**, and click **OK**:<br>• **Profile Name**—Enter a name to identify the profile.<br>• **Cutoff**—Specify a route withdrawal threshold above which a route advertisement is suppressed. Range: 0.0-1000.0. Default: 1.25.<br>• **Reuse**—Specify a route withdrawal threshold below which a suppressed route is used again. Range: 0.0-1000.0. Default: 5.<br>• **Max Hold Time (sec)**—Specify the maximum length of time in seconds that a route can be suppressed, regardless of how unstable it has been. Range: 0-3600 seconds. Default: 900 seconds.<br>• **Decay Half Life Reachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered reachable Range: 0-3600 seconds. Default: 300 seconds.<br>• **Decay Half Life Unreachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered unreachable. Range: 0 - 3600 seconds. Default: 300 seconds.<br>5. Click **OK**. |


© Palo Alto Networks, Inc.

| Configure BGP (Continued) | | |
|---|---|---|
| Step 5 | Configure the BGP peer group. | 1. Select the **Peer Group** tab and click **Add**.<br>2. Enter a **Name** for the peer group and select **Enable**.<br>3. Select the **Aggregated Confed AS Path** check box to include a path to the configured aggregated confederation AS.<br>4. Select the **Soft Reset with Stored Info** check box to perform a soft reset of the firewall after updating the peer settings.<br>5. Specify the type of peer or group from the **Type** drop-down and configure the associated settings (see below in this table for descriptions of Import Next Hop and Export Next Hop).<br>• **IBGP—Export Next Hop:** Specify **Original** or **Use self**<br>• **EBGP Confed—Export Next Hop:** Specify **Original** or **Use self**<br>• **EBGP Confed—Export Next Hop:** Specify **Original** or **Use self**<br>• **EBGP—Import Next Hop:** Specify **Original** or **Use self**, **Export Next Hop:** Specify **Resolve** or **Use self**. Select **Remove Private AS** if you want to force BGP to remove private AS numbers.<br>6. Click **OK** to save. |
| Step 6 | Configure Import and Export rules.<br>The import/export rules are used to import/export routes from/to other routers. For example, importing the default route from your Internet Service Provider. | 1. Select the **Import** tab and then click **Add** and enter a name in the **Rules** field and select the **Enable** check box.<br>2. Click **Add** and select the **Peer Group** to which the routes will be imported from.<br>3. Click the **Match** tab and define the options used to filter routing information. You can also define the Multi-Exit Discriminator (MED) value and a next hop value to routers or subnets for route filtering. The MED option is an external metric that lets neighbors know about the preferred path into an AS. A lower value is preferred over a higher value.<br>4. Click the **Action** tab and define the action that should occur (allow/deny) based on the filtering options defined in the **Match** tab. If **Deny** is selected, no further options need to be defined. If the **Allow** action is selected, define the other attributes.<br>5. Click the **Export** tab and define export attributes, which are similar to the **Import** settings, but are used to control route information that is exported from the firewall to neighbors.<br>6. Click **OK** to save. |

| Configure BGP (Continued) | | |
|---|---|---|
| Step 7 | Configure conditional advertising, which allows you to control what route to advertise in the event that a different route is not available in the local BGP routing table (LocRIB), indicating a peering or reachability failure. This is useful in cases where you want to try to force routes to one AS over another, for example if you have links to the Internet through multiple ISPs and you want traffic to be routed to one provider instead of the other unless there is a loss of connectivity to the preferred provider. | 1. Select the **Conditional Adv** tab, click **Add** and enter a name in the **Policy** field.<br>2. Select the **Enable** check box.<br>3. Click **Add** and in the **Used By** section enter the peer group(s) that will use the conditional advertisement policy.<br>4. Select the **Non Exist Filter** tab and define the network prefix(es) of the preferred route. This specifies the route that you want to advertise, if it is available in the local BGP routing table. If a prefix is going to be advertised and matches a Non Exist filter, the advertisement will be suppressed.<br>5. Select the **Advertise Filters** tab and define the prefix(es) of the route in the Local-RIB routing table that should be advertised in the event that the route in the non-exist filter is not available in the local routing table. If a prefix is going to be advertised and does not match a Non Exist filter, the advertisement will occur. |
| Step 8 | Configure aggregate options to summaries routes in the BGP configuration.<br><br>BGP route aggregation is used to control how BGP aggregates addresses. Each entry in the table results in one aggregate address being created. This will result in an aggregate entry in the routing table when at least one or more specific route matching the address specified is learned. | 1. Select the **Aggregate** tab, click **Add** and enter a name for the aggregate address.<br>2. In the **Prefix** field, enter the network prefix that will be the primary prefix for the aggregated prefixes.<br>3. Select the **Suppress Filters** tab and define the attributes that will cause the matched routes to be suppressed.<br>4. Select the **Advertise Filters** tab and define the attributes that will cause the matched routes to always be advertised to peers. |
| Step 9 | Configure redistribution rules.<br><br>This rule is used to redistribute host routes and unknown routes that are not on the local RIB to the peers routers. | 1. Select the **Redist Rules** tab and click **Add**.<br>2. In the **Name** field, enter an IP subnet or select a redistribution profile. You can also configure a new redistribution profile from the drop-down if needed.<br>3. Click the **Enable** check box to enable the rule.<br>4. In the **Metric** field, enter the route metric that will be used for the rule.<br>5. In the **Set Origin** drop-down, select **incomplete**, **igp**, or **egp**.<br>6. Optionally set MED, local preference, AS path limit and community values. |


© Palo Alto Networks, Inc.

# Session Settings and Timeouts

This section describes the global settings that affect TCP, UDP, and ICMPv6 sessions, in addition to IPv6, NAT64, NAT oversubscription, jumbo frame size, MTU, accelerated aging, and captive portal authentication. There is also a setting (Rematch Sessions) that allows you to apply newly configured security policies to sessions that are already in progress.

The **Device > Setup > Session** tab is where these session settings and timeouts are configured.

The first few topics below provide brief summaries of the Transport Layer of the OSI model, TCP, UDP, and ICMP. For more information about the protocols, refer to their respective RFCs. The remaining topics describe the session timeouts and settings.

- Transport Layer Sessions
- TCP
- UDP
- ICMP
- Configure Session Timeouts
- Configure Session Settings

## Transport Layer Sessions

A network session is a conversation that occurs between two or more communication devices, lasting for some period of time. A session is established and is torn down when the session ends. Different types of sessions occur at three layers of the OSI model: the Transport layer, the Session layer, and the Application layer.

The Transport Layer operates at Layer 4 of the OSI model, providing reliable or unreliable, end-to-end delivery and flow control of data. Internet protocols that implement sessions at the Transport layer include Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

## TCP

Transmission Control Protocol (TCP) (RFC 793) is one of the main protocols in the Internet Protocol (IP) suite, and is so prevalent that it is frequently referenced together with IP as *TCP/IP*. TCP is considered a reliable transport protocol because it provides error-checking while transmitting and receiving segments, it acknowledges segments received, and it reorders segments that arrive out of order from the device that sent them. TCP also requests and provides retransmission of segments that were dropped. TCP is stateful and connection-oriented, in that a connection between the sender and receiver is established for the duration of the session. TCP provides flow control of packets, so it can handle congestion over networks.

TCP performs a three-way handshake during session setup to initiate and acknowledge a session. After the data is transferred, the session is closed in an orderly manner.

Applications that use TCP as their transport protocol include Hypertext Transfer Protocol (HTTP), Hypertext Transfer Protocol Secure (HTTPS), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), Telnet, Post Office Protocol version 3 (POP3), Internet Message Access Protocol (IMAP), and Secure Shell (SSH).

The following topics describe details of the PAN-OS implementation of TCP.

- TCP Half Closed and TCP Time Wait Timers
- Unverified RST Timer

### TCP Half Closed and TCP Time Wait Timers

The TCP connection termination procedure uses a TCP Half Closed timer, which is triggered by the first FIN the firewall sees for a session. The timer is named TCP Half Closed because only one side of the connection has sent a FIN. A second timer, TCP Time Wait, is triggered by the second FIN or a RST.

If the firewall were to have only one timer triggered by the first FIN, a setting that was too short could prematurely close the half-closed sessions. Conversely, a setting that was too long would make the session table grow too much and possibly use up all of the sessions. Two timers allow you to have a relatively long TCP Half Closed timer and a short TCP Time Wait timer, thereby quickly aging fully closed sessions and controlling the size of the session table.

The following figure illustrates when the firewall's two timers are triggered during the TCP connection termination procedure.

 © Palo Alto Networks, Inc.

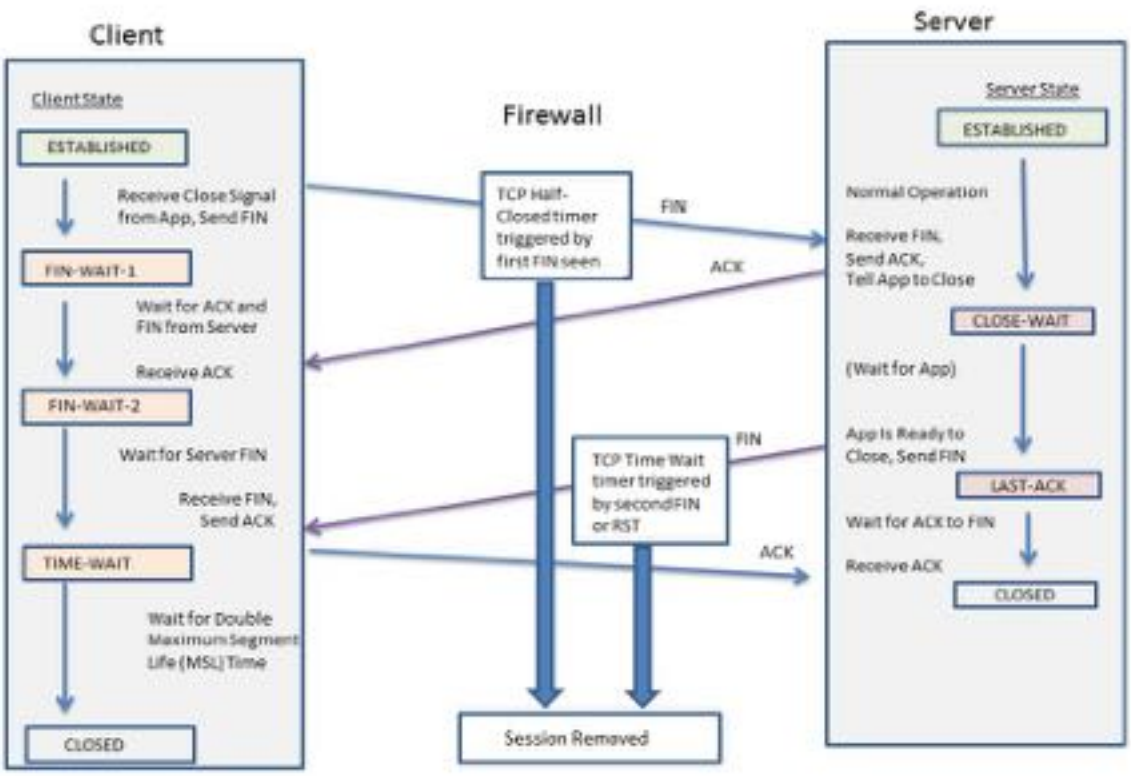


The TCP Time Wait timer should be set to a value less than the TCP Half Closed timer for the following reasons:

- The longer time allowed after the first FIN is seen gives the opposite side of the connection time to fully close the session.
- The shorter Time Wait time is because there is no need for the session to remain open for a long time after the second FIN or a RST is seen. A shorter Time Wait time frees up resources sooner, yet still allows time for the firewall to see the final ACK and possible retransmission of other datagrams.

If you configure a TCP Time Wait timer to a value greater than the TCP Half Closed timer, the commit will be accepted, but in practice the TCP Time Wait timer will not exceed the TCP Half Closed value.

The timers can be set globally or per application. The global settings are used for all applications by default. If you configure TCP wait timers at the application level, they override the global settings.

## Unverified RST Timer

If the firewall receives a Reset (RST) packet that cannot be verified (because it has an unexpected sequence number within the TCP window or it is from an asymmetric path), the Unverified RST timer controls the aging out of the session. It defaults to 30 seconds; the range is 1-600 seconds. The Unverified RST timer provides an additional security measure, explained in the second bullet below.

A RST packet will have one of three possible outcomes:

- A RST packet that falls outside the TCP window is dropped.

- A RST packet that falls inside the TCP window but does not have the exact expected sequence number is unverified and subject to the Unverified RST timer setting. This behavior helps prevent denial of service (DoS) attacks where the attack tries to disrupt existing sessions by sending random RST packets to the firewall.
- A RST packet that falls within the TCP window and has the exact expected sequence number is subject to the TCP Time Wait timer setting.

## UDP

User Datagram Protocol (UDP) (RFC 768) is another main protocol of the IP suite, and is an alternative to TCP. UDP is stateless and connectionless in that there is no handshake to set up a session, and no connection between the sender and receiver; the packets may take different routes to get to a single destination. UDP is considered an unreliable protocol because it does not provide acknowledgments, error-checking, retransmission, or reordering of datagrams. Without the overhead required to provide those features, UDP has reduced latency and is faster than TCP. UDP is referred to as a best-effort protocol because there is no mechanism or guarantee to ensure that the data will arrive at its destination.

Although UDP uses a checksum for data integrity, it performs no error checking at the network interface level. Error checking is assumed to be unnecessary or is performed by the application rather than UDP itself. UDP has no mechanism to handle flow control of packets.

UDP is often used for applications that require faster speeds and time-sensitive, real-time delivery, such as Voice over IP (VoIP), streaming audio and video, and online games. UDP is transaction-oriented, so it is also used for applications that respond to small queries from many clients, such as Domain Name System (DNS) and Trivial File Transfer Protocol (TFTP).

## ICMP

Internet Control Message Protocol (ICMP) (RFC 792) is another one of the main protocols of the Internet Protocol suite; it operates at the Network layer of the OSI model. ICMP is used for diagnostic and control purposes, to send error messages about IP operations, or messages about requested services or the reachability of a host or router. Network utilities such as traceroute and ping are implemented by using various ICMP messages.

ICMP is a connectionless protocol that does not open or maintain actual sessions. However, the ICMP messages between two devices can be considered a session.

Palo Alto Networks firewalls support ICMPv4 and ICMPv6. ICMPv4 and ICMPv6 error packets can be controlled by configuring a security policy for a zone, and selecting the **icmp** or **ipv6-icmp** application in the policy. Additionally, the ICMPv6 error packet rate can be controlled through the session settings, as described in the section Configure Session Settings.

 © Palo Alto Networks, Inc.

## ICMPv6 Rate Limiting

ICMPv6 rate limiting is a throttling mechanism to prevent flooding and DDoS attempts. The implementation employs an error packet rate and a token bucket, which work together to enable throttling and ensure that ICMP packets do not flood the network segments protected by the firewall.

First the global ICMPv6 error packet rate controls the rate at which ICMP error packets are allowed through the firewall; the default is 100 packets per second; the range is 10 to 65535 packets per second. If the firewall reaches the ICMP error packet rate, then the token bucket comes into play and throttling occurs, as follows.

The concept of a logical token bucket controls the rate at which ICMP messages can be transmitted. The number of tokens in the bucket is configurable, and each token represents an ICMP message that can be sent. The token count is decremented each time an ICMP message is sent; when the bucket reaches zero tokens, no more ICMP messages can be sent until another token is added to the bucket. The default size of the token bucket is 100 tokens (packets); the range is 10 to 65535 tokens.

To change the default token bucket size or error packet rate, see the section Configure Session Settings.

# Configure Session Timeouts

A session timeout defines the duration of time for which PAN-OS maintains a session on the firewall after inactivity in the session. By default, when the session timeout for the protocol expires, PAN-OS closes the session.

On the firewall, you can define a number of timeouts for TCP, UDP, and ICMP sessions in particular. The Default timeout applies to any other type of session. All of these timeouts are global, meaning they apply to all of the sessions of that type on the firewall.

In addition to the global settings, you have the flexibility to define timeouts for an individual application in the **Objects > Applications** tab. The firewall applies application timeouts to an application that is in established state. When configured, timeouts for an application override the global TCP or UDP session timeouts.

Returning to the global settings, perform the optional tasks below if you need to change default values of the global session timeout settings for TCP, UDP, ICMP, Captive Portal authentication, or other types of sessions. All values are in seconds.

The defaults are optimal values. However, you can modify these according to your network needs. Setting a value too low could cause sensitivity to minor network delays and could result in a failure to establish connections with the firewall. Setting a value too high could delay failure detection.

| **Change Session Timeouts** | | |
|---|---|---|
| Step 1 | Access the Session Settings. | 1. Select **Device > Setup > Session**.<br>2. In the Session Timeouts section, click the Edit icon. |

| Change Session Timeouts (Continued) | | |
|---|---|---|
| Step 2 | (Optional) Change miscellaneous timeouts. | • **Default**—Maximum length of time that a non-TCP/UDP or non-ICMP session can be open without a response. Default: 30. Range: 1-1599999.<br>• **Discard Default**—Maximum length of time that a non-TCP/UDP session remains open after PAN-OS denies a session based on security policies configured on the firewall. Default: 60. Range: 1-1599999.<br>• **Scan**—Maximum length of time that any session remains open after it is considered inactive; an application is regarded as inactive when it exceeds the application trickling threshold defined for the application. Default: 10. Range: 5-30.<br>• **Captive Portal**—Authentication session timeout for the Captive Portal web form. To access the requested content, the user must enter the authentication credentials in this form and be successfully authenticated. Default: 30. Range: 1-1599999.<br>To define other Captive Portal timeouts, such as the idle timer and the expiration time before the user must be re-authenticated, select **Device > User Identification > Captive Portal Settings**. See Configure Captive Portal. |
| Step 3 | (Optional) Change TCP timeouts. | • **Discard TCP**—Maximum length of time that a TCP session remains open after it is denied based on a security policy configured on the firewall. Default: 90. Range: 1-1599999.<br>• **TCP**—Maximum length of time that a TCP session remains open without a response, after a TCP session is in the Established state (after the handshake is complete and/or data is being transmitted). Default: 3600. Range: 1-1599999.<br>• **TCP Handshake**—Maximum length of time permitted between receiving the SYN-ACK and the subsequent ACK to fully establish the session. Default: 10. Range: 1-60.<br>• **TCP init**—Maximum length of time permitted between receiving the SYN and SYN-ACK prior to starting the TCP handshake timer. Default: 5. Range: 1-60.<br>• **TCP Half Closed**—Maximum length of time between receiving the first FIN and receiving the second FIN or a RST. Default: 120. Range: 1-604800.<br>• **TCP Time Wait**—Maximum length of time after receiving the second FIN or a RST. Default: 15. Range: 1-600.<br>• **Unverified RST**—Maximum length of time after receiving a RST that cannot be verified (the RST is within the TCP window but has an unexpected sequence number, or the RST is from an asymmetric path). Default: 30. Range: 1-600.<br>• See also the **Scan** timeout in the section (Optional) Change miscellaneous timeouts. |

 © Palo Alto Networks, Inc.

| Change Session Timeouts (Continued) | | |
|---|---|---|
| Step 4 | (Optional) Change UDP timeouts. | • **Discard UDP**—Maximum length of time that a UDP session remains open after it is denied based on a security policy configured on the firewall. Default: 60. Range: 1-1599999.<br>• **UDP**—Maximum length of time that a UDP session remains open without a UDP response. Default: 30. Range: 1-1599999.<br>• See also the **Scan** timeout in the section (Optional) Change miscellaneous timeouts. |
| Step 5 | (Optional) Change ICMP timeouts. | • **ICMP**—Maximum length of time that an ICMP session can be open without an ICMP response. Default: 6. Range: 1-1599999.<br>• See also the **Discard Default** and **Scan** timeout in the section (Optional) Change miscellaneous timeouts. |
| Step 6 | Commit the changes. | Click **OK** and **Commit** the changes. |

## Configure Session Settings

This topic describes various settings for sessions other than timeouts values. Perform these optional tasks if you need to change the default settings.

| Change Session Settings | | |
|---|---|---|
| Step 1 | Access the Session Settings. | 1. Select **Device > Setup > Session.**<br>2. In the **Session Settings** section, click the **Edit** icon. |

| Change Session Settings (Continued) | |
|---|---|
| Step 2 (Optional) Change session settings. | • **Rematch Sessions**—Causes the firewall to apply newly configured security policies to sessions that are already in progress. This capability is enabled by default. If this setting is disabled, any policy change applies only to sessions initiated after the policy change was committed.<br>For example, if a Telnet session started while an associated policy was configured that allowed Telnet, and you subsequently committed a policy change to deny Telnet, the firewall applies the revised policy to the current session and blocks it.<br>• **ICMPv6 Token Bucket Size**—Default: 100 tokens. See the section ICMPv6 Rate Limiting.<br>• **ICMPv6 Error Packet Rate (per sec)**—Default: 100. See the section ICMPv6 Rate Limiting.<br>• **Enable IPv6 Firewalling**—Enables firewall capabilities for IPv6. All IPv6-based configurations are ignored if IPv6 is not enabled. Even if IPv6 is enabled for an interface, the **IPv6 Firewalling** setting must also be enabled for IPv6 to function.<br>• **Enable Jumbo Frame**—Enables jumbo frame support on Ethernet interfaces. Jumbo frames have a maximum transmission unit (MTU) of 9216 bytes and are available on certain platforms.<br>**Global MTU**—Establishes the maximum transmission unit (MTU) available globally; this field is tied to the **Enable Jumbo Frame** field.<br>• If you do not check **Enable Jumbo Frame**, the **Global MTU** defaults to 1500 bytes; the range is 576 to 1500 bytes.<br>• If you check **Enable Jumbo Frame**, the **Global MTU** defaults to 9192 bytes; the range is 9192 to 9216 bytes.<br> If you enable jumbo frames and you have interfaces where the MTU is not specifically configured, those interfaces will automatically inherit the jumbo frame size. Therefore, before you enable jumbo frames, if you have any interface that you do not want to have jumbo frames, you must set the MTU for that interface to 1500 bytes or another value.<br>• **NAT64 IPv6 Minimum Network MTU**—Sets the global MTU for IPv6 translated traffic. The default of 1280 bytes is based on the standard minimum MTU for IPv6 traffic. |

 © Palo Alto Networks, Inc.

| Change Session Settings (Continued) | |
|---|---|
| | • **NAT Oversubscription Rate**—If NAT is configured to be Dynamic IP and Port (DIPP) translation, an oversubscription rate can be configured to multiply the number of times that the same translated IP address and port pair can be used concurrently. The rate is 1, 2, 4, or 8. The default setting is based on the firewall platform.<br>  • A rate of 1 means no oversubscription; each translated IP address and port pair can be used only once at a time.<br>  • If the setting is **Platform Default**, user configuration of the rate is disabled and the default oversubscription rate for the platform applies.<br>Reducing the oversubscription rate decreases the number of source device translations, but provides higher NAT rule capacities.<br>• **Accelerated Aging**—Enables faster aging-out of idle sessions. Select the check box to enable accelerated aging, and, if necessary, change the threshold (%) and scaling factor.<br>  • **Accelerated Aging Threshold**—Percentage of the session table that is full when accelerated aging begins. The default is 80%. When the session table reaches this threshold (% full), PAN-OS applies the **Accelerated Aging Scaling Factor** to the aging calculations for all sessions.<br>  • **Accelerated Aging Scaling Factor**—Scaling factor used in the accelerated aging calculations. The default scaling factor is 2, meaning that the accelerated aging occurs at a rate twice as fast as the configured idle time. The configured idle time divided by 2 results in a faster timeout of one-half the time. To calculate the session's accelerated aging, PAN-OS divides the configured idle time (for that type of session) by the scaling factor to determine a shorter timeout.<br>For example, if the scaling factor is 10, a session that would normally time out after 3600 seconds would time out 10 times faster (in 1/10 of the time), which is 360 seconds. |
| Step 3 Commit the changes. | Click **OK** and **Commit** the changes. |

# DHCP

This section describes Dynamic Host Configuration Protocol (DHCP) and the tasks required to configure an interface on a Palo Alto Networks firewall to act as a DHCP server, client, or relay agent. By assigning these roles to different interfaces, the firewall can perform multiple roles.


- DHCP Overview
- DHCP Messages
- DHCP Addressing
- DHCP Options
- Restrictions to Palo Alto Networks Implementation of DHCP
- Configure an Interface as a DHCP Server
- Configure an Interface as a DHCP Client
- Configure an Interface as a DHCP Relay Agent
- Monitor and Troubleshoot DHCP


## DHCP Overview

DHCP is a standardized protocol defined in RFC 2131, Dynamic Host Configuration Protocol. DHCP has two main purposes: to provide TCP/IP and link-layer configuration parameters and to provide network addresses to dynamically configured hosts on a TCP/IP network.

DHCP uses a client-server model of communication. This model consists of three roles that the device can fulfill: DHCP client, DHCP server, and DHCP relay agent.

- A device acting as a DHCP client (host) can request an IP address and other configuration settings from a DHCP server. Users on client devices save configuration time and effort, and need not know the network's addressing plan or other resources and options they are inheriting from the DHCP server.
- A device acting as a DHCP server can service clients. By using any of three DHCP Addressing mechanisms, the network administrator saves configuration time and has the benefit of reusing a limited number of IP addresses when a client no longer needs network connectivity. The server can deliver IP addressing and many DHCP options to many clients.
- A device acting as a DHCP relay agent transmits DHCP messages between DHCP clients and servers.

DHCP uses *User Datagram Protocol* (UDP), RFC 768, as its transport protocol. DHCP messages that a client sends to a server are sent to well-known port 67 (UDP—Bootstrap Protocol and DHCP). DHCP Messages that a server sends to a client are sent to port 68.

An interface on a Palo Alto Networks firewall can perform the role of a DHCP server, client, or relay agent. The interface of a DHCP server or relay agent must be a Layer 3 Ethernet, Aggregated Ethernet, or Layer 3 VLAN interface. The network administrator configures the firewall's interfaces with the appropriate settings for any combination of roles. The configuration tasks for each role are provided in this document.

 © Palo Alto Networks, Inc.

## DHCP Messages

DHCP uses eight standard message types, which are identified by an option type number in the DHCP message. For example, when a client wants to find a DHCP server, it broadcasts a DHCPDISCOVER message on its local physical subnetwork. If there is no DHCP server on its subnet and if DHCP Helper or DHCP Relay is configured properly, the message is forwarded to DHCP servers on a different physical subnet. Otherwise, the message will go no further than the subnet on which it originated. One or more DHCP servers will respond with a DHCPOFFER message that contains an available network address and other configuration parameters.

When the client needs an IP address, it sends a DHCPREQUEST to one or more servers. Of course if the client is requesting an IP address, it doesn't have one yet, so RFC 2131 requires that the broadcast message the client sends out have a source address of 0 in its IP header.

When a client requests configuration parameters from a server, it might receive responses from more than one server. Once a client has received its IP address, it is said that the client has at least an IP address and possibly other configuration parameters *bound* to it. DHCP servers manage such binding of configuration parameters to clients.

The following table lists the DHCP messages.

| DHCP Message | Description |
|---|---|
| DHCPDISCOVER | Client broadcast to find available DHCP servers. |
| DHCPOFFER | Server response to client's DHCPDISCOVER, offering configuration parameters. |
| DHCPREQUEST | Client message to one or more servers to do any of the following:<br>• Request parameters from one server and implicitly decline offers from other servers.<br>• Confirm that a previously allocated address is correct after, for example, a system reboot.<br>• Extend the lease of a network address. |
| DHCPACK | Server to client acknowledgment message containing configuration parameters, including a confirmed network address. |
| DHCPNAK | Server to client negative acknowledgment indicating the client's understanding of the network address is incorrect (for example, if the client has moved to a new subnet), or a client's lease has expired. |
| DHCPDECLINE | Client to server message indicating the network address is already being used. |
| DHCPRELEASE | Client to server message giving up the user of the network address and canceling the remaining time on the lease. |
| DHCPINFORM | Client to server message requesting only local configuration parameters; client has an externally configured network address. |

## DHCP Addressing

There are three ways that a DHCP server either assigns or sends an IP address to a client:

- **Automatic allocation**—The DHCP server assigns a permanent IP address to a client from its **IP Pools**. On the firewall, a **Lease** specified as **Unlimited** means the allocation is permanent.
- **Dynamic allocation**—The DHCP server assigns a reusable IP address from **IP Pools** of addresses to a client for a maximum period of time, known as a *lease*. This method of address allocation is useful when the customer has a limited number of IP addresses; they can be assigned to clients who need only temporary access to the network. See the DHCP Leases section.
- **Static allocation**—The network administrator chooses the IP address to assign to the client and the DHCP server sends it to the client. A static DHCP allocation is permanent; it is done by configuring a DHCP server and choosing a **Reserved Address** to correspond to the **MAC Address** of the client device. The DHCP assignment remains in place even if the client logs off, reboots, has a power outage, etc.

  Static allocation of an IP address is useful, for example, if you have a printer on a LAN and you do not want its IP address to keep changing, because it is associated with a printer name through DNS. Another example is if a client device is used for something crucial and must keep the same IP address, even if the device is turned off, unplugged, rebooted, or a power outage occurs, etc.

  Keep these points in mind when configuring a **Reserved Address**:

  - It is an address from the **IP Pools**. You may configure multiple reserved addresses.
  - If you configure no **Reserved Address**, the clients of the server will receive new DHCP assignments from the pool when their leases expire or if they reboot, etc. (unless you specified that a **Lease** is **Unlimited**).
  - If you allocate all of the addresses in the **IP Pools** as a **Reserved Address,** there are no dynamic addresses free to assign to the next DHCP client requesting an address.
  - You may configure a **Reserved Address** without configuring a **MAC Address**. In this case, the DHCP server will not assign the **Reserved Address** to any device. You might reserve a few addresses from the pool and statically assign them to a fax and printer, for example, without using DHCP.

## DHCP Leases

A lease is defined as the time period for which a DHCP server allocates a network address to a client. The lease might be extended (renewed) upon subsequent requests. If the client no longer needs the address, it can release the address back to the server before the lease is up. The server is then free to assign that address to a different client if it has run out of unassigned addresses.

The lease period configured for a DHCP server applies to all of the addresses that a single DHCP server (interface) dynamically assigns to its clients. That is, all of that interface's addresses assigned dynamically are of **Unlimited** duration or have the same **Timeout** value. A different DHCP server configured on the firewall may have a different lease term for its clients. A **Reserved Address** is a static address allocation and is not subject to the lease terms.

Per the DHCP standard, RFC 2131, a DHCP client does not wait for its lease to expire, because it risks getting a new address assigned to it. Instead, when a DHCP client reaches the halfway point of its lease period, it attempts to extend its lease so that it retains the same IP address. Thus, the lease duration is like a sliding window.

Typically if an IP address was assigned to a device, the device was subsequently taken off the network and its lease was not extended, the DHCP server will let that lease run out. Because the client is gone from the network and no longer needs the address, the lease duration in the server is reached and the lease is in "Expired" state.

 © Palo Alto Networks, Inc.

The firewall has a hold timer that prevents the expired IP address from being reassigned immediately. This behavior temporarily reserves the address for the device in case it comes back onto the network. But if the address pool runs out of addresses, the server re-allocates this expired address before the hold timer expires. Expired addresses are cleared automatically as the systems needs more addresses or when the hold timer releases them.

In the CLI, use the `show dhcp server lease` operational command to view lease information about the allocated IP addresses. If you do not want to wait for expired leases to be released automatically, you can use the `clear dhcp lease interface value expired-only` command to clear expired leases, making those addresses available in the pool again. You can use the `clear dhcp lease interface value ip ip` command to release a particular IP address. Use the `clear dhcp lease interface value mac mac_address` command to release a particular MAC address.

## DHCP Options

The history of DHCP and DHCP options traces back to the Bootstrap Protocol (BOOTP). BOOTP was used by a host to configure itself dynamically during its booting procedure. A host could receive an IP address and a file from which to download a boot program from a server, along with the server's address and the address of an Internet gateway.

Included in the BOOTP packet was a vendor information field, which could contain a number of tagged fields containing various types of information, such as the subnet mask, the BOOTP file size, and many other values. RFC 1497 describes the *BOOTP Vendor Information Extensions*.

These extensions eventually expanded with the use of DHCP and DHCP host configuration parameters, also known as options. Similar to vendor extensions, DHCP options are tagged data items that provide information to a DHCP client. The options are sent in a variable-length field at the end of a DHCP message. For example, the DHCP Message Type is option 53, and a value of 1 indicates the DHCPDISCOVER message. DHCP options are defined in RFC 2132, DHCP Options and BOOTP Vendor Extensions.

A DHCP client can negotiate with the server, limiting the server to send only those options that the client requests.

The Palo Alto Networks firewall supports a subset of the standard DHCP options for its implementation of DHCP server. Such options are configured on the DHCP server and sent to the clients that sent a DHCPREQUEST to the server. The clients are said to *inherit* and implement the options that they are programmed to accept. The firewall supports the following predefined options on its DHCP servers, shown in the order in which they appear on the **DHCP Server** configuration screen:

| DHCP Option | DHCP Option Name |
|---|---|
| 51 | Lease duration |
| 3 | Gateway |
| 1 | IP Pool Subnet (mask) |
| 6 | Domain Name System (DNS) server address (primary and secondary) |
| 44 | Windows Internet Name Service (WINS) server address (primary and secondary) |

| DHCP Option | DHCP Option Name |
|---|---|
| 41 | Network Information Service (NIS) server address (primary and secondary) |
| 42 | Network Time Protocol (NTP) server address (primary and secondary) |
| 70 | Post Office Protocol Version 3 (POP3) server address |
| 69 | Simple Mail Transfer Protocol (SMTP) server address |
| 15 | DNS suffix |

## Restrictions to Palo Alto Networks Implementation of DHCP

- The DHCP standard, RFC 2131, is designed to support IPv4 and IPv6 addresses. The Palo Alto Networks implementations of DHCP server and DHCP client support IPv4 addresses only. Its DHCP relay implementation supports IPv4 and IPv6.
- DHCP client is not supported in High Availability active/active mode.
- The firewall supports DHCPv4 Server and DHCPv6 Relay. However, a single interface cannot support both DHCPv4 Server and DHCPv6 Relay.
- DHCP replaces BOOTP; BOOTP is not supported

## Configure an Interface as a DHCP Server

The prerequisites for this task are:

- ❒ You have configured a Layer 3 Ethernet or Layer 3 VLAN interface.
- ❒ The interface is assigned to a virtual router and a zone.
- ❒ You know a valid pool of IP addresses from your network plan that can be designated to be assigned by your DHCP server to clients.

Perform the following task to configure an interface on the firewall to act as a DHCP server. You may configure multiple DHCP servers.

 © Palo Alto Networks, Inc.

| Configure an Interface as a DHCP Server | |
|---|---|
| Step 1 Select an interface to be a DHCP Server. | 1. Select **Network > DHCP > DHCP Server** and click **Add**.<br>2. Enter an **Interface** name or select one from the drop-down.<br>3. For **Mode**, select **enabled** or **auto** mode. Auto mode enables the server and disables it if another DHCP server is detected on the network. The **disabled** setting disables the server.<br>4. Optionally click the **Ping IP when allocating new IP** check box if you want the server to ping the IP address before it assigns that address to its client.<br> If the ping receives a response, that means a different device already has that address, so it is not available. The server assigns the next address from the pool instead. This behavior is similar to Optimistic Duplicate Address Detection (DAD) for IPv6, RFC 4429.<br> After setting options and returning to the DHCP server tab, the **Probe IP** column for the interface indicates if this check box was checked. |

| Configure an Interface as a DHCP Server (Continued) | |
|---|---|
| Step 2 Configure the DHCP Options that the server sends to its clients. | • In the Options section, select a **Lease** type:<br>  • **Unlimited** causes the server to dynamically choose IP addresses from the **IP Pools** and assign them permanently to clients.<br>  • **Timeout** determines how long the lease will last. Enter the number of **Days** and **Hours**, and optionally the number of **Minutes**.<br>• **Inheritance Source**—Leave **None** or select a source DHCP client interface or PPPoE client interface to propagate various server settings into the DHCP server. If you specify an **Inheritance Source**, select one or more options below that you want **inherited** from this source.<br>Benefits of specifying an inheritance source are that it quickly adds DHCP options from the upstream server of the source DHCP client. It also keeps the client's options updated if the source changes an option. For example, if the source replaces its NTP server (which had been identified as the **Primary NTP** server), the client will automatically inherit the new address as its **Primary NTP** server.<br>• **Check inheritance source status**—If your selected an **Inheritance Source**, clicking this link opens the **Dynamic IP Interface Status** window, which displays the options that were inherited from the DHCP client.<br>• **Gateway**—IP address of the network gateway (an interface on the firewall) that is used to reach any device not on the same LAN as this DHCP server.<br>• **Ippool Subnet**—Network mask that applies to the addresses in the **IP Pools**. |


© Palo Alto Networks, Inc.

| Configure an Interface as a DHCP Server (Continued) | | |
|---|---|---|
| | | For the following fields, click the down arrow and select **None**, or **inherited**, or enter a remote server's IP address that your DHCP server will send to clients for accessing that service. If you select **inherited,** the DHCP server inherits the values from the source DHCP client specified as the **Inheritance Source**.<br>The DHCP server sends these settings to its clients.<br>• **Primary DNS, Secondary DNS**—IP address of the preferred and alternate Domain Name System (DNS) servers.<br>• **Primary WINS, Secondary WINS**—IP address of the preferred and alternate Windows Internet Naming Service (WINS) servers.<br>• **Primary NIS, Secondary NIS**—IP address of the preferred and alternate Network Information Service (NIS) servers.<br>• **Primary NTP, Secondary NTP**—IP address of the available Network Time Protocol servers.<br>• **POP3 Server**—IP address of a Post Office Protocol (POP3) server.<br>• **SMTP Server**—IP address of a Simple Mail Transfer Protocol (SMTP) server.<br>• **DNS Suffix**—Suffix for the client to use locally when an unqualified hostname is entered that it cannot resolve. |
| Step 3 | Identify the stateful pool of IP addresses from which the DHCP server chooses an address and assigns it to a DHCP client.<br> If you are not the network administrator for your network, ask the network administrator for a valid pool of IP addresses from the network plan that can be designated to be assigned by your DHCP server. | 1. In the **IP Pools** field, click **Add** and enter the range of IP addresses from which this server assigns an address to a client. Enter an IP subnet and subnet mask (for example, 192.168.1.0/24) or a range of IP addresses (for example, 192.168.1.10-192.168.1.20). At least one IP pool is required.<br>2. Optionally repeat Step 1 to specify another IP address pool. |
| Step 4 | (Optional) Specify an IP address from the IP pools that will not be assigned dynamically. If you also specify a **MAC Address**, the **Reserved Address** is assigned to that device when the device requests an IP address through DHCP.<br> See the DHCP Addressing section for an explanation of allocation of a **Reserved Address**. | 1. In the **Reserved Address** field, click **Add.**<br>2. Enter an IP address from the **IP Pools** (format *x.x.x.x*) that you do not want to be assigned dynamically by the DHCP server.<br>3. Optionally specify the **MAC Address** (format *xx:xx:xx:xx:xx:xx*) of the device to which you want to permanently assign the IP address specified in Step 2.<br>4. Optionally repeat Steps 2-3 to reserve another address. |
| Step 5 | Save the configuration. | Click **OK** and **Commit** the change. |

## Configure an Interface as a DHCP Client

Before configuring a firewall interface as a DHCP client, make sure you have configured a Layer 3 Ethernet or Layer 3 VLAN interface, and the interface is assigned to a virtual router and a zone. Perform this task if you need to use DHCP to request an IPv4 address for an interface on your firewall.

| **Configure an Interface as a DHCP Client** | | |
|---|---|---|
| Step 1 | Configure an interface as a DHCP client. | 1. Select **Network** > **Interfaces**.<br>2. On the **Ethernet** tab or the **VLAN** tab, click **Add** and enter an interface, or click a configured interface, that you want to be a DHCP client.<br>The following shows the **Ethernet Interface** screen:<br> <br>3. Click the **IPv4** tab; for **Type**, select **DHCP Client**.<br>4. Select the **Enable** check box.<br>5. Optionally select the **Automatically create default route pointing to default gateway provided by server** check box. This causes the firewall to create a static route to a default gateway that will be useful when clients are trying to access many destinations that do not need to have routes maintained in a routing table on the firewall.<br>6. Optionally enter a **Default Route Metric** (priority level) for the route between the firewall and the DHCP server. A route with a lower number has higher priority during route selection. For example, a route with a metric of 10 is used before a route with a metric of 100. Range is 1-65535. There is no default metric.<br>7. Optionally click **Show DHCP Client Runtime Info** to see all of the settings the client has inherited from its DHCP server. |
| Step 2 | Save the configuration. | Click **OK** and **Commit** the change.<br>Now the Ethernet interface indicates **Dynamic-DHCP Client** in its **IP Address** field on the **Ethernet** tab. |
| Step 3 | (Optional) See which interfaces on the firewall are configured as DHCP clients. | 1. Select **Network > Interfaces > Ethernet** and look in the **IP Address** field to see which interfaces indicate DHCP Client.<br>2. Select **Network > Interfaces > VLAN** and look in the **IP Address** field to see which interfaces indicate DHCP Client. |


© Palo Alto Networks, Inc.

## Configure an Interface as a DHCP Relay Agent

Before configuring a firewall interface as a DHCP relay agent, make sure you have configured a Layer 3 Ethernet or Layer 3 VLAN interface, and the interface is assigned to a virtual router and a zone. You want that interface be able to pass DHCP messages between clients and servers. The interface can forward messages to a maximum of four external DHCP servers. A client DHCPDISCOVER message is sent to all configured servers, and the DHCPOFFER message of the first server that responds is relayed back to the requesting client.

Perform the following task to configure an interface on the firewall to act as a DHCP relay agent:

| **Configure an Interface as a DHCP Relay Agent** | | |
|---|---|---|
| Step 1 | Select DHCP Relay. | Select **Network** > **DHCP > DHCP Relay**. |
| Step 2 | Specify the IP address of each DHCP server with which the DHCP relay agent will communicate. | 1. In the **Interface** field, select from the drop-down the interface you want to be the DHCP relay agent.<br>2. Check either the **IPv4** or **IPv6** check box, indicating the type of DHCP server address you will specify.<br>3. If you checked **IPv4**, in the **DHCP Server IP Address** field, click **Add**. Enter the address of the DHCP server to and from which you will relay DHCP messages.<br>4. If you checked **IPv6**, in the **DHCP Server IPv6 Address** field, click **Add**. Enter the address of the DHCP server to and from which you will relay DHCP messages. If you specify a *multicast* address, also specify an outgoing **Interface**.<br>5. Optionally repeat Steps 2-4 to enter a maximum of four DHCP server addresses. |
| Step 3 | Save the configuration. | Click **OK** and **Commit** the change. |

## Monitor and Troubleshoot DHCP

You can view the status of dynamic address leases that your DHCP server has assigned or that your DHCP client has been assigned by issuing commands from the CLI. You can also clear leases before they time out and are released automatically. For complete command syntax, see the *PAN-OS Command Line Interface Reference Guide*.

- View DHCP Server Information
- Clear Leases Before They Expire Automatically
- View DHCP Client Information
- Gather Debug Output about DHCP

## View DHCP Server Information

To view DHCP pool statistics, IP addresses the server has assigned, the corresponding MAC address, state and duration of the lease, and time the lease began, use the following command. If the address was configured as a **Reserved Address**, the `state` column indicates `reserved` and there is no `duration` or `lease_time`. If the lease was configured as **Unlimited**, the `duration` column displays a value of `0`.

```
admin@PA-200> show dhcp server lease all
interface: "ethernet1/2"
Allocated IPs: 1, Total number of IPs in pool: 5. 20.0000% used
ip              mac               state      duration    lease_time
192.168.3.11    f0:2f:af:42:70:cf  committed  0           Wed Jul  2 08:10:56 2014
admin@PA-200>
```

To view the options that a DHCP server has assigned to clients, use the following command:

```
admin@PA-200> show dhcp server settings all
Interface       GW            DNS1           DNS2        DNS-Suffix    Inherit source
------------------------------------------------------------------------------------
ethernet1/2     192.168.3.1   10.43.2.10     10.44.2.10                ethernet1/3
admin@PA-200>
```

## Clear Leases Before They Expire Automatically

The following example shows how to release expired DHCP Leases of an interface (server) before the hold timer releases them automatically. Those addresses will be available in the IP pool again.

```
admin@PA-200> clear dhcp lease interface ethernet1/2 expired-only
```

The following example shows how to release the lease of a particular IP address:

```
admin@PA-200> clear dhcp lease interface ethernet1/2 ip 192.168.3.1
```

The following example shows how to release the lease of a particular MAC address:

```
admin@PA-200> clear dhcp lease interface ethernet1/2 mac f0:2c:ae:29:71:34
```

## View DHCP Client Information

To view the status of IP address leases sent to the firewall when it is acting as a DHCP client, use the `show dhcp client state` *`interface_name`* command or the following command:

```
admin@PA-200> show dhcp client state all
Interface       State        IP              Gateway         Leased-until
--------------------------------------------------------------------------
ethernet1/1     Bound        10.43.14.80     10.43.14.1      70315
admin@PA-200>
```

## Gather Debug Output about DHCP

To gather debug output about DHCP, use one of the following commands:

```
admin@PA-200> debug dhcpd
admin@PA-200> debug management-server dhcpd
```

 © Palo Alto Networks, Inc.

# NAT

This section describes Network Address Translation (NAT) and how to configure NAT rules and features.


- Purpose of NAT
- NAT Rules and Security Policies
- Source NAT and Destination NAT
- NAT Rule Capacities
- Dynamic IP and Port NAT Oversubscription
- Dataplane NAT Memory Statistics
- Configure NAT


## Purpose of NAT

NAT was introduced to solve the problem of an organization not having enough public, globally-routable IPv4 addresses assigned to it by the Internet Assigned Numbers Authority (IANA) or a Regional Internet Registry. NAT translates private, non-routable IPv4 addresses to one or more globally-routable IPv4 addresses, thereby conserving an organization's routable IP addresses.

Since its origination, the use of NAT has expanded. For example, NAT is used as a way to not disclose the real IP addresses of hosts that need access to public addresses. It is also used to manage traffic by performing port forwarding. The NAT64 option translates between IPv6 and IPv4 addresses, providing connectivity between networks using disparate IP addressing schemes, and therefore a migration path to IPv6 addressing. PAN-OS supports all of these uses.

If you use private IP addresses within your internal networks, you must use NAT to translate the private addresses to public addresses that can be routed on external networks. In PAN-OS, you create NAT policy rules that instruct the firewall which packet addresses need translation and what the translated addresses are.

## NAT Rules and Security Policies

You configure a NAT rule to match a packet's source zone and destination zone, at a minimum. In addition to zones, you can configure matching criteria based on the packet's destination interface, source and destination address, and service. You can configure multiple NAT rules. The firewall evaluates the rules in order from the top down. Once a packet matches the criteria of a single NAT rule, the packet is not subjected to additional NAT rules. Therefore, your list of NAT rules should be in order from most specific to least specific so that packets are subjected to the most specific rule you created for them.

It is important to understand the firewall's flow logic when it applies NAT rules and security policies so that you can determine what rules you need, based on the zones you have defined.

© Palo Alto Networks, Inc.

Upon ingress, the firewall inspects the packet and does a route lookup to determine the egress interface and zone. Then the firewall determines if the packet matches one of the NAT rules that have been defined, based on source and/or destination zone. It then evaluates and applies any security policies that match the packet based on the original (pre-NAT) source and destination addresses, but the post-NAT zones. Finally, upon egress, for a matching NAT rule, the firewall translates the source and/or destination address and port numbers.

Keep in mind that the translation of the IP address and port do not occur until the packet leaves the firewall. The NAT rules and security policies apply to the original IP address (the pre-NAT address). A NAT rule is configured based on the zone associated with a pre-NAT IP address.

Security policies differ from NAT rules because security policies examine post-NAT zones to determine whether the packet is allowed or not. Because the very nature of NAT is to modify source or destination IP addresses, which can result in modifying the packet's outgoing interface and zone, security policies are enforced on the post-NAT zone.

### Address Pools Identified as Address Objects

When configuring a NAT IP address pool, it is typical to identify it as address object. It can be a host IP address, IP address range, or IP subnet. Because both NAT rules and security policies use address objects, it is a best practice to distinguish between them by naming an address object used for NAT with a prefix, such as "NAT-*name*."

## Source NAT and Destination NAT

The firewall supports both source address and/or port translation and destination address and/or port translation.

- Source NAT
- Destination NAT

### Source NAT

Source NAT is typically used by internal users to access the Internet; the source address is translated and thereby kept private. There are three types of source NAT:

- **Dynamic IP and Port (DIPP)**—Allows multiple hosts to have their source IP addresses translated to the same public IP address with different port numbers. The dynamic translation is to the next available address in the NAT address pool, which you configure as a **Translated Address** pool be to an IP address, range of addresses, a subnet, or a combination of these.

  As an alternative to the next address in the NAT address pool, DIPP allows you to specify the address of the **Interface** itself. The advantage of specifying the interface in the NAT rule is that the NAT rule will be automatically updated to use any address subsequently acquired by the interface.

  DIPP has a default NAT oversubscription rate, which is the number of times that the same translated IP address and port pair can be used concurrently. For more information, see Dynamic IP and Port NAT Oversubscription and Modify the Oversubscription Rate for DIPP NAT.

 © Palo Alto Networks, Inc.

- **Dynamic IP**—Allows the 1-to-1, dynamic translation of a source IP address only (no port number) to the next available address in the NAT address pool. By default, if the source address pool is larger than the NAT address pool and eventually all of the NAT addresses are allocated, new connections that need address translation are dropped. The default behavior can be changed by clicking **Advanced (Dynamic IP/Port Fallback)**, which causes DIPP addresses to be used when necessary. In either event, as sessions terminate and the addresses in the pool become available, they can be allocated to translate new connections.
- **Static IP**—Allows the 1-to-1, static translation of a source IP address, but leaves the source port unchanged. A common scenario for a static IP translation is an internal server that must be available to the Internet.

### Destination NAT

Destination NAT is performed on incoming packets, when the firewall translates a public destination address to a private address. Destination NAT does not use address pools or ranges. It is a 1-to-1, static translation with the option to perform port forwarding or port translation.

- **Static IP**—Allows the 1-to-1, static translation of a destination IP address and optionally the port number.

One common use of destination NAT is to configure several NAT rules that map a single public destination address to several private destination addresses assigned to servers or services. For example:

  - **Port Forwarding**—Can translate a public destination address and port number to a private destination address, but keeps the same port number.
  - **Port Translation**—Can translate a public destination address and port number to a private destination address and a different port number, thus keeping the real port number private. It is configured by entering a **Translated Port** on the **Translated Packet** tab in the NAT policy rule.

## NAT Rule Capacities

The number of NAT rules allowed is based on the firewall platform. Individual rule limits are set for static, Dynamic IP (DIP), and Dynamic IP and Port (DIPP) NAT. The sum of the number of rules used for these NAT types cannot exceed the total NAT rule capacity. For DIPP, the rule limit is based on the device's oversubscription setting (8, 4, 2, or 1) and the assumption of one translated IP address per rule. To see platform-specific NAT rule limits and translated IP address limits, use the Compare Firewalls tool.

Consider the following when working with NAT rules:

- If you run out of pool resources, you cannot create more NAT rules, even if the platform's maximum rule count has not been reached.
- If you consolidate NAT rules, the logging and reporting will also be consolidated. The statistics are provided per the rule, not per all of the addresses within the rule. If you need granular logging and reporting, do not combine the rules.

© Palo Alto Networks, Inc.

## Dynamic IP and Port NAT Oversubscription

Dynamic IP and Port (DIPP) NAT allows you to use each translated IP address and port pair multiple times (8, 4, or 2 times) in concurrent sessions. This reusability of an IP address and port (known as oversubscription) provides scalability for customers who have too few public IP addresses. The design is based on the assumption that hosts are connecting to different destinations, therefore sessions can be uniquely identified and collisions are unlikely. The oversubscription rate in effect multiplies the original size of the address/port pool to 8, 4, or 2 times the size. For example, the default limit of 64K concurrent sessions allowed, when multiplied by an oversubscription rate of 8, results in 512K concurrent sessions allowed.

The oversubscription rates that are allowed vary based on the platform. The oversubscription rate is global; it applies to the device. This oversubscription rate is set by default and consumes memory, even if you have enough public IP addresses available to make oversubscription unnecessary. You can reduce the rate from the default setting to a lower setting or even 1 (which means no oversubscription). By configuring a reduced rate, you decrease the number of source device translations possible, but increase the DIP and DIPP NAT rule capacities. To change the default rate, see Modify the Oversubscription Rate for DIPP NAT.

If you select **Platform Default**, your explicit configuration of oversubscription is turned off and the default oversubscription rate for the platform applies, as shown in the table below. The **Platform Default** setting allows for an upgrade or downgrade of a software release.

The following table lists the default (highest) oversubscription rate for each platform.

| Platform | Default Oversubscription Rate |
|---|---|
| PA-200 | 2 |
| PA-500 | 2 |
| PA-2020 | 2 |
| PA-2050 | 2 |
| PA-3020 | 2 |
| PA-3050 | 2 |
| PA-3060 | 2 |
| PA-4020 | 4 |
| PA-4050 | 8 |
| PA-4060 | 8 |
| PA-5020 | 4 |
| PA-5050 | 8 |
| PA-5060 | 8 |
| PA-7050 | 8 |
| VM-100 | 1 |
| VM-200 | 1 |

 © Palo Alto Networks, Inc.

| Platform | Default Oversubscription Rate |
|---|---|
| VM-300 | 2 |
| VM-1000-HV | 2 |

The firewall supports a maximum of 256 translated IP addresses per NAT rule, and each platform supports a maximum number of translated IP addresses (for all NAT rules combined). If oversubscription causes the maximum translated addresses per rule (256) to be exceeded, the firewall will automatically reduce the oversubscription ratio in an effort to have the commit succeed. However, if your NAT rules result in translations that exceed the maximum translated addresses for the platform, the commit will fail.

## Dataplane NAT Memory Statistics

The `show running global-ippool` command displays statistics related to NAT memory consumption for a pool. The Size column displays the number of bytes of memory that the resource pool is using. The Ratio column displays the oversubscription ratio (for DIPP pools only). The lines of pool and memory statistics are explained in the following sample output:

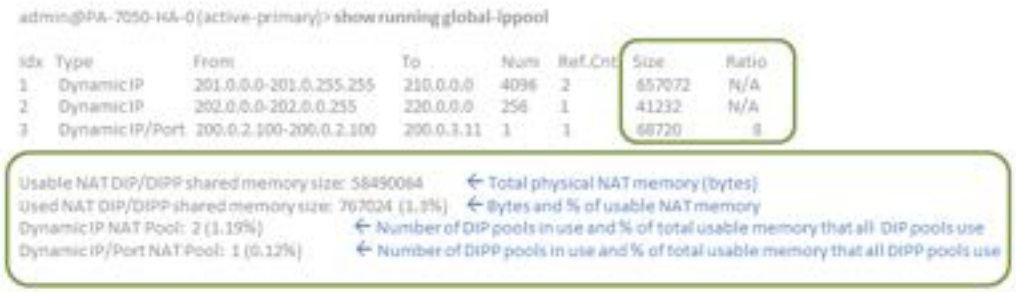


For NAT pool statistics for a virtual system, the `show running ippool` command has columns indicating the memory size used per NAT rule and the oversubscription ratio used (for DIPP rules). The following is sample output for the command.

A field in the output of the `show running nat-rule-ippool rule` command shows the memory (bytes) used per NAT rule. The following is sample output for the command, with the memory usage for the rule encircled.

 © Palo Alto Networks, Inc.

## Configure NAT

Perform the following tasks to configure various aspects of NAT.

- Translate Internal Client IP Addresses to Your Public IP Address
- Enable Clients on the Internal Network to Access your Public Servers
- Enable Bi-Directional Address Translation for Your Public-Facing Servers
- Modify the Oversubscription Rate for DIPP NAT
- Disable NAT for a Specific Host or Interface

The NAT example in this section is based on the following topology, which was also used in Getting Started for setting up interfaces and zones:

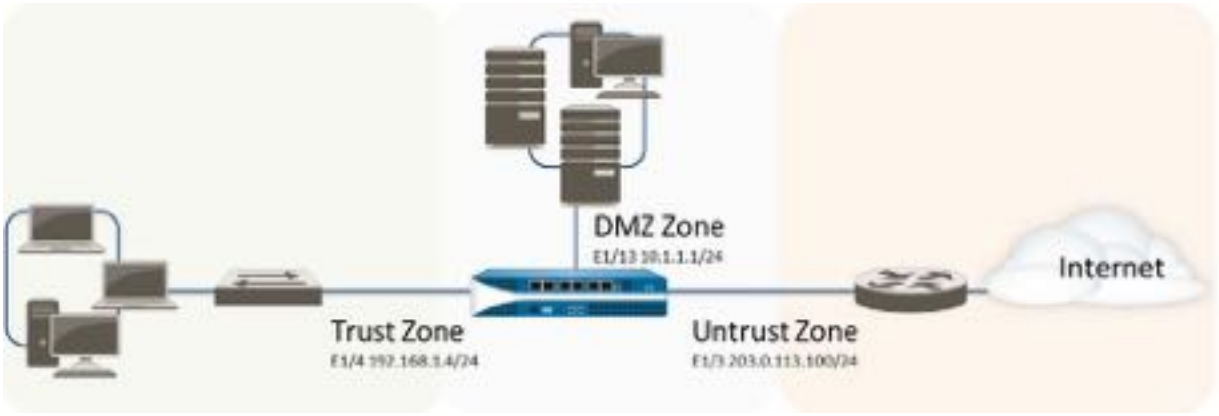


Based on the topology initially used in Getting Started to create the interfaces and zones, there are three NAT policies we need to create as follows:

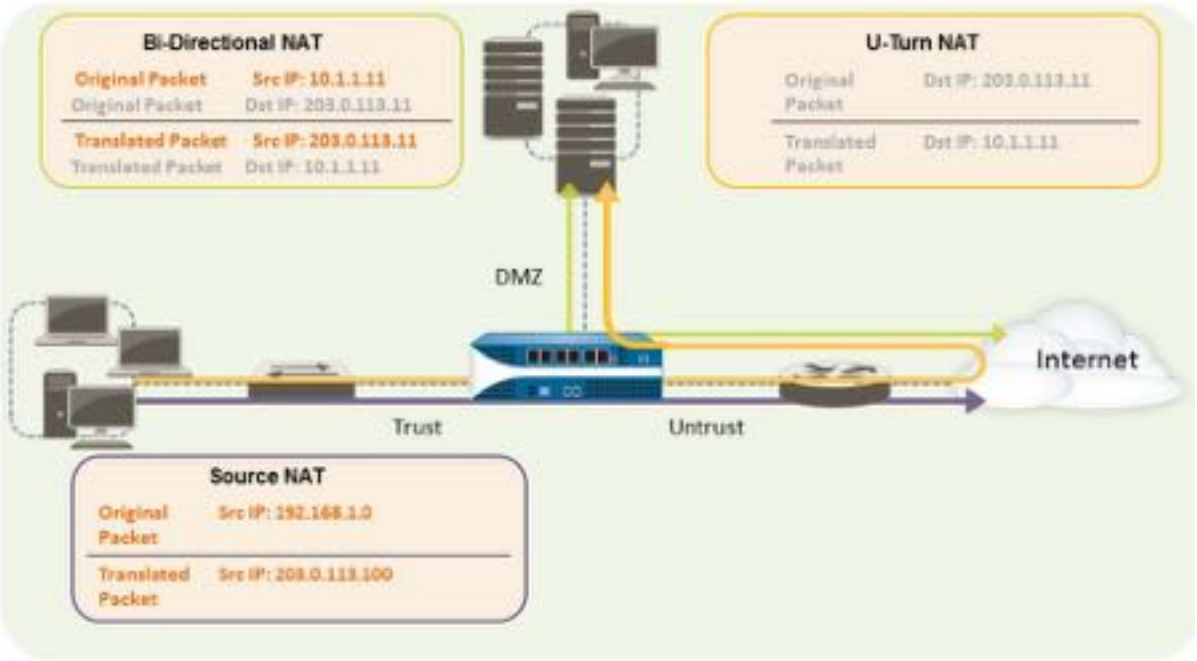


- To enable the clients on the internal network to access resources on the Internet, the internal 192.168.1.0 addresses will need to be translated to publicly routable addresses. In this case, we will configure source NAT, using the egress interface address, 203.0.113.100, as the source address in all packets that leave the firewall from the internal zone. See Translate Internal Client IP Addresses to Your Public IP Address for instructions.
- To enable clients on the internal network to access the public web server in the DMZ zone, we will need to configure a NAT rule that redirects the packet from the external network, where the original routing table lookup will determine it should go based on the destination address of 203.0.113.11 within the packet, to the actual address of the web server on the DMZ network of 10.1.1.11. To do this you must create a NAT rule from the trust zone (where the source address in the packet is) to the untrust zone (where the original destination address is) to translate the destination address to an address in the DMZ zone. This type of destination NAT is called *U-Turn NAT*. See Enable Clients on the Internal Network to Access your Public Servers for instructions.
- To enable the web server—which has both a private IP address on the DMZ network and a public-facing address for access by external users—to both send and receive requests, the firewall must translate the incoming packets from the public IP address to the private IP address and the outgoing packets from the private IP address to the public IP address. On the firewall, you can accomplish this with a single bi-directional static source NAT policy. See Enable Bi-Directional Address Translation for Your Public-Facing Servers.

 © Palo Alto Networks, Inc.

## Translate Internal Client IP Addresses to Your Public IP Address

When a client on your internal network sends a request, the source address in the packet contains the IP address for the client on your internal network. If you use private IP address ranges internally, the packets from the client will not be able to be routed on the Internet unless you translate the source IP address in the packets leaving the network into a publicly routable address.

On the firewall you can do this by configuring a source NAT policy that translates the source address and optionally the port into a public address. One way to do this is to translate the source address for all packets to the egress interface on your firewall, as shown in the following procedure.

| **Configure Source NAT** | |
|---|---|
| Step 1 Create an address object for the external IP address you plan to use. | 1. Select **Objects > Addresses** and then click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the object.<br>3. Select **IP Netmask** from the **Type** drop-down and then enter the IP address of the external interface on the firewall, 203.0.113.100 in this example.<br> <br>4. To save the address object, click **OK**.<br>Although you do not have to use address objects in your policies, it is a best practice because it simplifies administration by allowing you to make updates in one place rather than having to update every policy where the address is referenced. |

| Configure Source NAT (Continued) | |
|---|---|
| Step 2 Create the NAT policy. | 1. Select **Policies > NAT** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the policy.<br>3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.<br> <br>4. On the **Translated Packet** tab, select **Dynamic IP And Port** from the **Translation Type** drop-down in the **Source Address Translation** section of the screen and then click **Add**.<br>5. Select the address object you just created.<br> 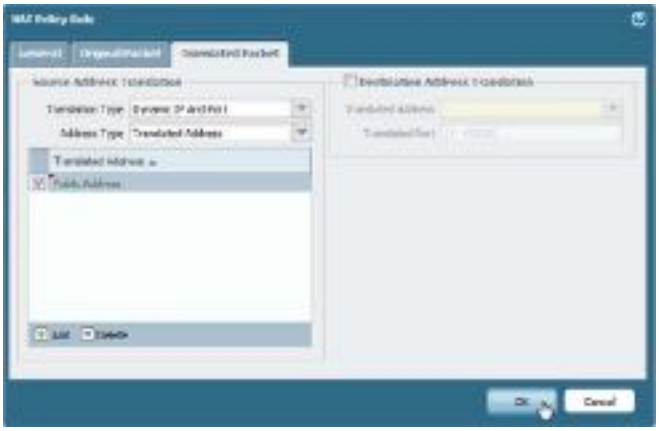 <br>6. Click **OK** to save the NAT policy. |
| Step 3 Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

## Enable Clients on the Internal Network to Access your Public Servers

When a user on the internal network sends a request for access to the corporate web server in the DMZ, the DNS server will resolve to the public IP address. When processing the request, the firewall will use the original destination in the packet (the public IP address) and route the packet to the egress interface for the untrust zone. In order for the firewall to know that it must translate the public IP address of the web server to an address on the DMZ network when it receives requests from users on the trust zone, you must create a destination NAT rule that will enable the firewall to send the request to the egress interface for the DMZ zone as follows.

| **Configure U-Turn NAT** | |
|---|---|
| Step 1 Create an address object for the web server. | 1. Select **Objects > Addresses** and click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the object.<br>3. Select **IP Netmask** from the **Type** drop-down and enter the public IP address of the web server, 203.0.113.11 in this example.<br> <br>4. Click **OK**. |
| Step 1 Create the NAT policy. | 1. Select **Policies > NAT** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the NAT rule.<br>3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.<br>4. In the **Destination Address** section, click **Add** and select the address object you created for your public web server.<br> |

| Configure U-Turn NAT (Continued) | |
|---|---|
| | 5. On the **Translated Packet** tab, select the **Destination Address Translation** check box and then enter the IP address that is assigned to the web server interface on the DMZ network, 10.1.1.11 in this example.<br> <br>6. Click **OK** to save the NAT policy. |
| Step 2 Save the configuration. | Click **Commit**. |

## Enable Bi-Directional Address Translation for Your Public-Facing Servers

When your public-facing servers have private IP addresses assigned on the network segment where they are physically located, you need a source NAT rule to translate the source address of the server to the external address upon egress. You create a static NAT rule to translate the internal source address, 10.1.1.11, to the external web server address, 203.0.113.11 in our example.

However, a public-facing server must be able to both send and receive packets. You need a reciprocal policy that translates the public address (the destination IP address in incoming packets from Internet users) into the private address so that the firewall can route the packet to your DMZ network. You create a bi-directional static NAT rule, as described in the following procedure.

| Configure Bi-Directional NAT | |
|---|---|
| Step 1 Create an address object for the web server's internal IP address. | 1. Select **Objects > Addresses** and click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the object.<br>3. Select **IP Netmask** from the **Type** drop-down and enter the IP address of the web server on the DMZ network, 10.1.1.11 in this example.<br> <br>4. Click **OK**.<br>If you did not already create an address object for the *public* address of your web server, you should create that object now. |

 © Palo Alto Networks, Inc.

| Configure Bi-Directional NAT (Continued) | |
|---|---|
| Step 2 Create the NAT policy. | 1. Select **Policies > NAT** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the NAT rule.<br>3. On the **Original Packet** tab, select the zone you created for your DMZ in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.<br>4. In the **Source Address** section, click **Add** and select the address object you created for your internal web server address.<br> 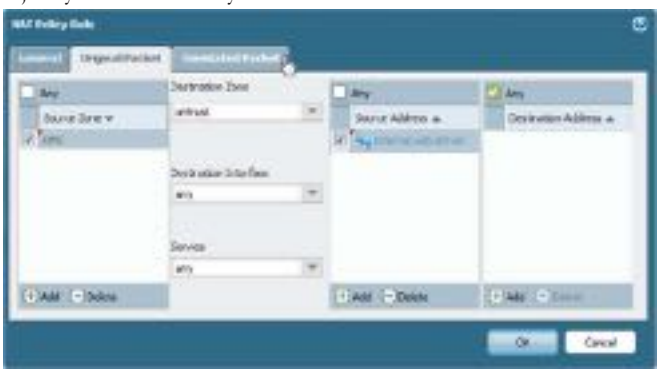 <br>5. On the **Translated Packet** tab, select **Static IP** from the **Translation Type** drop-down in the **Source Address Translation** section and then select the address object you created for your external web server address from the **Translated Address** drop-down.<br>6. In the **Bi-directional** field, select **Yes**.<br> <br>7. Click **OK** to save the NAT policy. |
| Step 3 Save the configuration. | Click **Commit**. |

## Modify the Oversubscription Rate for DIPP NAT

If you have enough public IP addresses that you do not need to use DIPP NAT oversubscription, you can reduce the oversubscription rate and thereby gain more DIP and DIPP NAT rules allowed.

| **Set NAT Oversubscription** | | |
|---|---|---|
| Step 1 | View the DIPP NAT oversubscription rate. | 1. Select **Device > Setup > Session > Session Settings.** View the **NAT Oversubscription Rate** setting. |
| Step 2 | Set the DIPP NAT oversubscription rate. | 1. Click the Edit icon in the Session Settings section.<br>2. In the **NAT Oversubscription Rate** drop-down, select **1x**, **2x**, **4x**, or **8x**, depending on which ratio you want.<br>The **Platform Default** setting applies the default oversubscription setting for the platform. If you want no oversubscription, select **1x**.<br>3. Click **OK** and **Commit** the change. |

## Disable NAT for a Specific Host or Interface

You may have exceptions where you do not want NAT to occur for a certain host in a subnet or for traffic exiting a specific interface. Source NAT and destination NAT rules can be configured to disable translation in those cases. The following procedure shows how to disable source NAT for a host.

| **Create a Source NAT Exemption** | | |
|---|---|---|
| Step 1 | Create the NAT policy. | 1. Select **Policies > NAT** and click **Add**.<br>2. Enter a descriptive **Name** for the policy.<br>3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.<br>4. For **Source Address**, click **Add** and enter the host address. Click **OK**.<br>5. On the **Translated Packet** tab, select **None** from the **Translation Type** drop-down in the **Source Address Translation** section of the screen.<br>6. Click **OK** to save the NAT policy. |
| Step 2 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

# LACP

The firewall can now use Link Aggregation Control Protocol (LACP) to detect the physical interfaces between itself and a connected device (peer) and manage those interfaces as a single virtual interface (aggregate group). An aggregate group increases the bandwidth between peers. Enabling LACP provides redundancy within the group: the protocol automatically detects interface failures and performs failover to standby interfaces. Without LACP, you must manually identify interface failures occurring at layers above the physical or occurring between peers that are not directly connected.

The following Palo Alto Networks firewalls support LACP: PA-500, PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7050. The firewalls support LACP for HA3 (only on the PA-500, PA-3000 Series, PA-4000 Series, and PA-5000 Series), Layer 2, and Layer 3 interfaces.

The following topics describe LACP and how to configure it for aggregate groups:

- LACP Settings
- Configure LACP

## LACP Settings

To implement LACP, you must configure an aggregate group on each LACP peer that will use the protocol. Before you Configure LACP, determine the optimal settings for each peer, as described in the following topics:

- Mode
- Transmission Rate
- Fast Failover
- Port Priority and System Priority

### Mode

In LACP active mode, a device actively searches the network for peers and queries their status (available or unresponsive). In passive mode, a device only responds to queries from an active peer. Between any two peers, it is recommended that one be active and the other passive. LACP cannot function if both peers are passive.

### Transmission Rate

You can configure a device to detect the state (available or unresponsive) of peers and individual interfaces at fast (every second) or slow (every 30 seconds) intervals. When the device does not receive an LACP update from the peer within a period that is three times the transmission rate (3 seconds or 90 seconds), it flags the peer as unresponsive. Therefore, set the transmission rate according to how much LACP processing your network can support and how quickly a device should detect and resolve interface failures.

© Palo Alto Networks, Inc.

For example, in a network where peers have many redundant links, a single interface failure would not disrupt traffic. In this case, a slow rate would suffice for detecting failures and would avoid the extra processing associated with the fast rate. If the network can support the extra processing and requires high availability (for example, datacenters supporting critical business operations), a fast transmission rate would enable the firewall to substitute failed interfaces quicker.

If devices have different transmission rates, each uses the rate of its peer.

## Fast Failover

When an interface goes down, it fails over to a standby interface. The IEEE 802.1ax standard that defines LACP specifies a failover process that takes at least three seconds. For networks that require quicker failover, the firewall provides an option that performs fast failover within one second. This option is recommended for deployments in which critical data might be lost during the standard failover interval. For example, if the amount of traffic between peers is close to the maximum that the active interfaces can support and one interface fails, the risk of data being lost while a standby interface becomes active is significantly less for fast failover than for standard failover.

Because both peers perform failover when an interface goes down, it is recommended that you configure fast failover on both so that they complete the process within the same time frame.

## Port Priority and System Priority

When you configure an aggregate group, the **Max Ports** parameter determines how many interfaces can be active at any given time. If the number of interfaces you assign to the group exceeds the **Max Ports**, the remaining interfaces are in standby mode. If an active interface fails, a standby interface becomes active. LACP uses the **LACP Port Priority** assigned to each interface to determine which interfaces are initially active and to determine the order in which standby interfaces become active upon failover. (You set port priorities when configuring individual interfaces, not when configuring the aggregate group.)

It is recommended that you assign a unique priority number to each interface in an aggregate group. However, if multiple interfaces have the same priority number, the interface with the lowest port number has the highest priority.

To enable LACP between two peers, each requires an aggregate group configuration. If the port priority values for the member interfaces differ in each peer, the values of the peer with the higher **System Priority** will override the other peer in determining active and standby interfaces.

For port and system priorities, the priority level is the inverse of its numerical value: a priority of 1 designates the highest priority level; a value of 65535 designates the lowest priority level. The default is 32768.

The following figure shows an aggregate group instance configured on a firewall and on a switch. The group has four interfaces and the **Max Ports** parameter is set to three. That means three interfaces are active and one is standby. The firewall administrator assigned port priorities 12, 64, 128, and 258 to the interfaces. The switch administrator assigned port priorities 22, 54, 158, and 228 to the interfaces. LACP uses the port priorities of the firewall instance because the **System Priority** of the firewall is higher than that of the switch.

 © Palo Alto Networks, Inc.

**Figure: LACP System Priority**

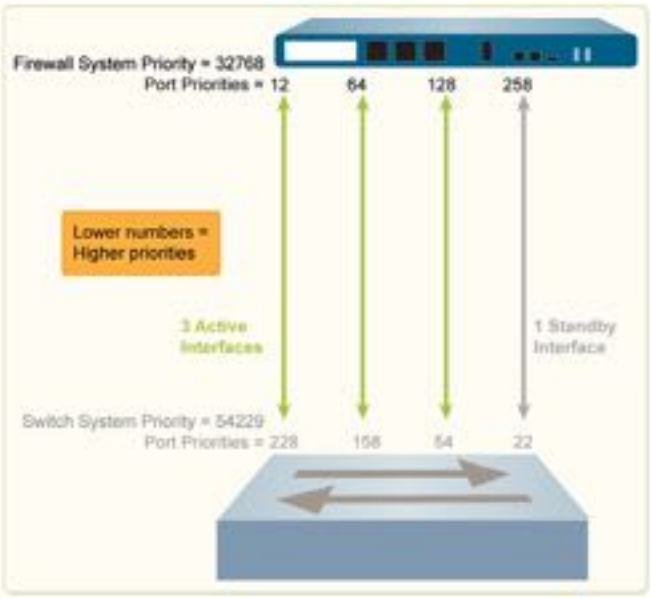


It is a best practice to assign a unique system priority to each peer. However, in some cases, peers might have the same value. This could happen if a different administrator configured the aggregate group on each peer. In such cases, the peer for which the *system ID* is a lower numerical value will override the other peer with respect to port priorities. LACP automatically derives the system ID from the system priority and system MAC address. A numerically lower MAC address produces a numerically lower system ID.

The following figure shows the same peers as in Figure: LACP System Priority but with identical system priorities. In this case, LACP uses the system IDs to determine port prioritization: the switch overrides the firewall.

**Figure: LACP System ID**

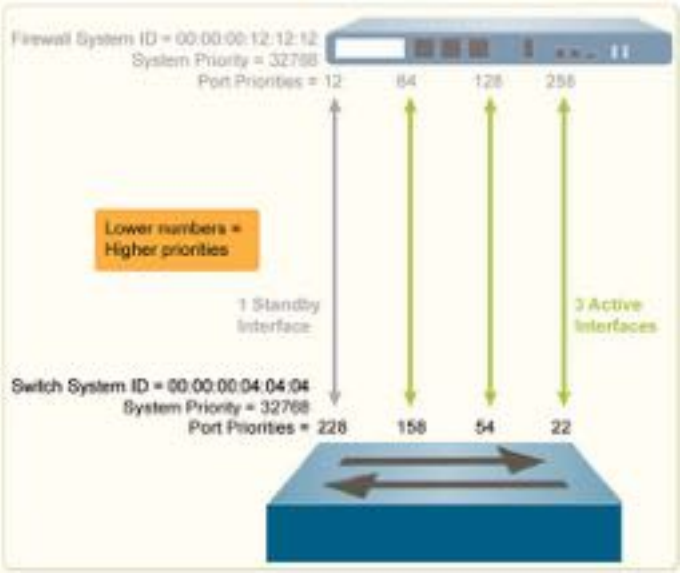

Firewalls in a high availability (HA) pair have the same system priority value. However, in an active/passive deployment, the system ID for each can be the same or different, depending on whether you assign the same MAC address. (Firewalls in an active/active deployment require unique MAC addresses so PAN-OS automatically assigns them.) When the LACP peers (also in HA mode) are virtualized (appearing to the network as a single device), selecting the **Same System MAC Address for Active-Passive HA** option for the firewalls is a best practice to minimize latency during failover. When the LACP peers are not virtualized, using the unique MAC address of each firewall is the best practice to minimize failover latency. In the latter case, if the HA firewall pair and the LACP peer devices all have the same system priority but each firewall in the HA pair has a unique system ID, one firewall might have a lower system ID than the LACP peers while the other firewall has a higher system ID than the LACP peers. In this case, when failover occurs on the firewalls, port prioritization switches between the LACP peers and the firewall that becomes active.

## Configure LACP

Before starting this procedure:

- Determine which physical interfaces connect the LACP peers. This procedure assumes the cabling is complete.
- Determine the optimal LACP Settings for the peers.

This procedure only covers the configuration steps for an LACP peer that is a Palo Alto Networks firewall. Other devices have proprietary methods for configuring LACP.

Perform the following steps to configure LACP on a firewall. In a high availability deployment, configure the primary (for active/active) or active (for active/passive) firewall; the secondary or passive firewall automatically synchronizes with it.

 © Palo Alto Networks, Inc.

| Configure LACP | |
|---|---|
| Step 1 Add an aggregate group with LACP enabled. | 1. Select **Network > Interfaces > Ethernet** and click **Add Aggregate Group**.<br>2. In the field adjacent to the read-only **Interface Name**, enter a number (1-8) to identify the group.<br>3. For the **Interface Type**, select **HA**, **Layer 2**, or **Layer 3**.<br>Only select **HA** if the interface is an HA3 link between two firewalls in an active/active deployment.<br>4. On the **LACP** tab, select **Enable LACP**.<br>5. For the LACP **Mode**, select whether the firewall is **Passive** (default) or **Active**.<br>6. For the **Transmission Rate**, select **Fast** or **Slow**.<br>7. If desired, select **Fast Failover** (under one second) for when interfaces go down. Otherwise, the standard failover (at least three seconds) applies.<br>8. Enter a **System Priority** number (1-65535, default 32768) that determines whether the firewall or peer overrides the other with respect to port priorities. Note that the lower the number, the higher the priority.<br>9. For the **Max Ports**, enter the number of interfaces (1-8) that are active. The value cannot exceed the number of interfaces you assign to the group. If the number of assigned interfaces exceeds the number of active interfaces, the remaining interfaces will be in standby mode.<br>10. By default, each firewall in an HA pair has a unique MAC address. A **Same System MAC Address for Active-Passive HA** configuration is recommended only when the LACP peers are virtualized (appearing to the network as a single device). In such cases, select the check box and select the system-generated **MAC Address**, or enter your own. You must verify the address is globally unique.<br>If the firewalls are not in active/passive HA mode, PAN-OS ignores your selections for these fields. Firewalls in an active/active deployment require unique MAC addresses so PAN-OS automatically assigns them. If the **Interface Type** is **HA**, these fields do not appear.<br>11. Click **OK**. |

| Configure LACP (Continued) | | |
|---|---|---|
| Step 2 | Assign interfaces to the aggregate group. | Perform the following steps for each physical interface (1-8) that will belong to the aggregate group.<br> During normal operations, assign more than one interface. You would only assign a single interface if a troubleshooting procedure requires it.<br>1. Select **Network > Interfaces** and select the interface name. It must be the same type (HA, Layer 2, or Layer 3) as the aggregate group.<br>2. Change the **Interface Type** to **Aggregate Ethernet**.<br>3. Select the **Aggregate Group** you just defined.<br>4. Select the **Link Speed**, **Link Duplex**, and **Link State**. It is a best practice to set the same link speed and duplex values for every interface in the group. For non-matching values, the commit operation displays a warning and PAN-OS defaults to the higher speed and full duplex.<br>5. Enter an **LACP Port Priority** (1-65535, default 32768). If the number of interfaces you assign exceeds the **Max Ports** value you configured for the group, the port priorities determine which interfaces are active or standby. The lower the numeric value, the higher the priority.<br>6. Click **OK** and **Commit**. |
| Step 3 | Verify the status of the aggregate group. | 1. Select **Network > Interfaces > Ethernet**.<br>2. Verify that the Features column displays an LACP enabled icon for the aggregate group.<br>3. Verify that the Link State column displays a green icon for the aggregate group, indicating that all member interfaces are up. If the icon is yellow, at least one member is down but not all. If the icon is red, all members are down.<br>4. If the Link State icon for the aggregate group is yellow or red, select **Monitor > Logs > System** and review the logs for the LACP subtype to investigate the causes of the interface failures. |
| Step 4 | (Optional) If you selected **HA** as the **Interface Type** for the aggregate group, enable forwarding of packets over the HA3 link. | 1. Select **Device > High Availability > Active/Active Config** and edit the Packet Forwarding section.<br>2. For the **HA3 Interface**, select the aggregate group you configured.<br>3. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# Policy

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Captive Portal, Denial of Service (DoS), and Zone protection policies. All these different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. The following topics describe how to work with policy:

- Policy Types
- Security Policy
- Policy Objects
- Security Profiles
- Enumeration of Rules Within a Rulebase
- Use Tags to Visually Distinguish Objects
- Tips and Tricks for Searching for Objects in Security Policy
- Use a Dynamic Block List in Policy
- Register IP Addresses and Tags Dynamically
- Monitor Changes in the Virtual Environment
- CLI Commands for Dynamic IP Addresses and Tags
- Policy-Based Forwarding
- DoS Protection Against Flooding of New Sessions

# Policy Types

The Palo Alto Networks next-generation firewall supports a variety of policy types that work together to safely enable applications on your network.

| Policy Type | Description |
|---|---|
| Security | Determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service. For more details, see Security Policy. |
| NAT | Instruct the firewall which packets need translation and how to do the translation. The firewall supports both source address and/or port translation and destination address and/or port translation. For more details, see NAT. |
| QoS | Identify traffic requiring QoS treatment (either preferential treatment or bandwidth-limiting) using a defined parameter or multiple parameters and assigns it a class. For more details, see Quality of Service. |
| Policy Based Forwarding | Identify traffic that should use a different egress interface than the one that would normally be used based on the routing table. For details, see Policy-Based Forwarding. |
| Decryption | Identify encrypted traffic that you want to inspect for visibility, control, and granular security. For more details, see Decryption. |
| Application Override | Identify sessions that you do not want processed by the App-ID engine, which is a Layer-7 inspection. Traffic matching an application override policy forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4. For more details, see Manage Custom or Unknown Applications. |
| Captive Portal | Identify traffic that requires the user to be known. The captive portal policy is only triggered if other User-ID mechanisms did not identify a user to associate with the source IP address. For more details, see Captive Portal. |
| DoS Protection | Identify potential denial-of-service (DoS) attacks and take protective action in response to rule matches. DoS Protection Profile. |

 © Palo Alto Networks, Inc.

# Security Policy

Security policies protect network assets from threats and disruptions and aid in optimally allocating network resources for enhancing productivity and efficiency in business processes. On the Palo Alto Networks firewall, security policies determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service.

All traffic passing through the firewall is matched against a session and each session is matched against a security policy. When a session match occurs, the security policy is applied to bi-directional traffic (client to server and server to client) in that session. For traffic that doesn't match any defined rules, the default rules apply. The default rules—displayed at the bottom of the security rulebase—are predefined to allow all intrazone (within the zone) traffic and deny all interzone (between zones) traffic. Although these rules are part of the pre-defined configuration and are read-only by default, you can override them and change a limited number of settings, including the tags, action (allow or deny), log settings, and security profiles.

Security policies are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log, if logging is enabled for that rule. The logging options are configurable for each rule, and can for example be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.


- Components of a Security Policy Rule
- Security Policy Best Practices

# Components of a Security Policy Rule

The security policy rule construct permits a combination of the required and optional fields as detailed in the following tables:

- Required Fields
- Optional Fields

## Required Fields

| Required Field | Description |
|---|---|
| **Name** | A label that supports up to 31 characters, used to identify the rule. |
| **Rule Type** | Specifies whether the rule applies to traffic within a zone, between zones, or both:<br>• **universal** (default)—Applies the rule to all matching interzone and intrazone traffic in the specified source and destination zones. For example, if you create a universal role with source zones A and B and destination zones A and B, the rule would apply to all traffic within zone A, all traffic within zone B, and all traffic from zone A to zone B and all traffic from zone B to zone A.<br>• **intrazone**—Applies the rule to all matching traffic within the specified source zones (you cannot specify a destination zone for intrazone rules). For example, if you set the source zone to A and B, the rule would apply to all traffic within zone A and all traffic within zone B, but not to traffic between zones A and B.<br>• **interzone**—Applies the rule to all matching traffic between the specified source and destination zones. For example, if you set the source zone to A, B, and C and the destination zone to A and B, the rule would apply to traffic from zone A to zone B, from zone B to zone A, from zone C to zone A, and from zone C to zone B, but not traffic within zones A, B, or C. |
| **Source Zone** | The zone from which the traffic originates. |
| **Destination Zone** | The zone at which the traffic terminates. If you use NAT, make sure to always reference the post-NAT zone. |
| **Application** | The application which you wish to control. The firewall uses App-ID, the traffic classification technology, to identify traffic on your network. App-ID provides application control and visibility in creating security policies that block unknown applications, while enabling, inspecting, and shaping those that are allowed. |
| **Action** | Specifies an *Allow* or *Deny* action for the traffic based on the criteria you define in the rule. |

 © Palo Alto Networks, Inc.

## Optional Fields

| Optional Field | Description |
|---|---|
| **Tag** | A keyword or phrase that allows you to filter security rules. This is handy when you have defined many rules and wish to then review those that are tagged with a particular keyword, for example *Inbound to DMZ*. |
| **Description** | A text field, up to 255 characters, used to describe the rule. |
| **Source IP Address** | Define host IP or FQDN, subnet, named groups, or country-based enforcement. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address). |
| **Destination IP Address** | The location or destination for the traffic. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address). |
| **User** | The user or group of users for whom the policy applies. You must have User-ID enabled on the zone. To enable User-ID, see User-ID Overview. |
| **URL Category** | Using the URL Category as match criteria allows you to customize security profiles (antivirus, anti-spyware, vulnerability, file-blocking, Data Filtering, and DoS) on a per-URL-category basis. For example, you can prevent.exe file download/upload for URL categories that represent higher risk while allowing them for other categories. This functionality also allows you to attach schedules to specific URL categories (allow social-media websites during lunch & after-hours), mark certain URL categories with QoS (financial, medical, and business), and select different log forwarding profiles on a per-URL-category-basis.<br>Although you can manually configure URL categories on your device, to take advantage of the dynamic URL categorization updates available on the Palo Alto Networks firewalls, you must purchase a URL filtering license.<br>To block or allow traffic based on URL category, you must apply a URL Filtering profile to the security policy rules. Define the URL Category as *Any* and attach a URL Filtering profile to the security policy. See Set Up Basic Security Policies for information on using the default profiles in your security policy and see Control Access to Web Content for more details. |
| **Service** | Allows you to select a Layer 4 (TCP or UDP) port for the application. You can choose *any*, specify a port, or use *application-default* to permit use of the standards-based port for the application. For example, for applications with well- known port numbers such as DNS, the *application-default* option will match against DNS traffic only on TCP port 53. You can also add a custom application and define the ports that the application can use.<br>For inbound allow rules (for example, from untrust to trust), using *application-default* prevents applications from running on unusual ports and protocols. Application-default is the default option; while the device still checks for all applications on all ports, with this configuration, applications are only allowed on their standard ports/protocols. |
| **Security Profiles** | Provide additional protection from threats, vulnerabilities, and data leaks. Security profiles are only evaluated for rules that have an *allow* action. |

| Optional Field | Description (Continued) |
|---|---|
| **HIP Profile** (for GlobalProtect) | Allows you to identify clients with Host Information Profile (HIP) and then enforce access privileges. |
| **Options** | Allow you to define logging for the session, log forwarding settings, change Quality of Service (QoS) markings for packets that match the rule, and schedule when (day and time) the security rule should be in effect. |

 © Palo Alto Networks, Inc.

## Security Policy Best Practices

The task of safely enabling Internet access and preventing misuse of web access privileges, and exposure to vulnerabilities and attacks is a continuous process. The key principle when defining policy on the Palo Alto Networks firewall is to use a positive enforcement approach. Positive enforcement implies that you selectively allow what is required for day-to-day business operations as opposed to a negative enforcement approach where you would selectively block everything that is not allowed. Consider the following suggestions when creating policy:

- If you have two or more zones with identical security requirements, combine them into one security rule.
- The ordering of rules is crucial to ensure the best match criteria. Because policy is evaluated top down, the more specific policy must precede the ones that are more general, so that the more specific rule is not *shadowed*. The term shadow refers to a rule that is not evaluated or is skipped because it is placed lower in the policy list. When the rule is placed lower, it is not evaluated because the match criteria was met by another rule that preceded it, thereby shadowing the rule from policy evaluation.
- To restrict and control access to inbound applications, in the security policy, explicitly define the port that the service/application will be listening on.
- Logging for broad allow rules—for example access to well known servers like DNS—can generate a lot of traffic. Hence it is not recommended unless absolutely necessary.
- By default, the firewall creates a log entry at the end of a session. However, you can modify this default behavior and configure the firewall to log at the start of the session. Because this significantly increases the log volume, logging at session start is recommended only when you are troubleshooting an issue. Another alternative for troubleshooting without enabling logging at session start is to use the session browser (**Monitor > Session Browser**) to view the sessions in real time.

© Palo Alto Networks, Inc.

# Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With policy objects that are a collective unit, you can reference the object in security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

You can create the following policy objects on the firewall:

| Policy Object | Description |
|---|---|
| Address/Address Group, Region | Allow you to group specific source or destination addresses that require the same policy enforcement. The address object can include an IPv4 or IPv6 address (single IP, range, subnet) or the FQDN. Alternatively, a region can be defined by the latitude and longitude coordinates or you can select a country and define an IP address or IP range. You can then group a collection of address objects to create an *address group* object.<br>You can also use dynamic address groups to dynamically update IP addresses in environments where host IP addresses change frequently. |
| User/User Group | Allow you to create a list of users from the local database or an external database and group them. |
| Application Group and Application Filter | An *Application Filter* allows you to filter applications dynamically. It allows you to filter, and save a group of applications using the attributes defined in the application database on the firewall. For example, you can filter by one or more attributes—category, sub-category, technology, risk, characteristics—and save your application filter. With an application filter, when a PAN-OS content update occurs, any new applications that match your filter criteria are automatically added to your saved application filter.<br>An *Application Group* allows you to create a static group of specific applications that you wish to group together for a group of users or for a particular service. |
| Service/Service Groups | Allows you to specify the source and destination ports and protocol that a service can use. The firewall includes two pre-defined services—service-http and service-https—that use TCP ports 80 and 8080 for HTTP, and TCP port 443 for HTTPS. You can however, create any custom service on any TCP/UDP port of your choice to restrict application usage to specific ports on your network (in other words, you can define the default port for the application).<br> To view the standard ports used by an application, in **Objects > Applications** search for the application and click the link. A succinct description displays. |

 © Palo Alto Networks, Inc.

# Security Profiles

While security policies enable you to allow or deny traffic on your network, security profiles help you define an *allow but scan* rule, which scan allowed applications for threats, such as viruses, malware, spyware, and DDOS attacks. When traffic matches the allow rule defined in the security policy, the security profile(s) that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.

Security profiles are not used in the match criteria of a traffic flow. The security profile is applied to scan traffic after the application or category is allowed by the security policy.

The firewall provides default security profiles that you can use out of the box to begin protecting your network from threats. See Set Up Basic Security Policies for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, you can create custom profiles. See Scan Traffic for Threats for more information.

You can add security profiles that are commonly applied together to a security profile group; this set of profiles can be treated as a unit and added to security policies in one step (or included in security policies by default, if you choose to set up a default security profile group).

The following topics provide more detailed information about each type of security profile and how to set up a security profile group:

- Antivirus Profiles
- Anti-spyware Profiles
- Vulnerability Protection Profiles
- URL Filtering Profiles
- Data Filtering Profiles
- File Blocking Profiles
- DoS Protection Profile
- Zone Protection Profiles
- Security Profile Group

© Palo Alto Networks, Inc.

## Antivirus Profiles

Antivirus profiles protect against viruses, worms, and trojans as well as spyware downloads. Using a stream-based malware prevention engine, which inspects traffic the moment the first packet is received, the Palo Alto Networks antivirus solution can provide protection for clients without significantly impacting the performance of the firewall. This profile scans for a wide variety of malware in executables, PDF files, HTML and JavaScript viruses, including support for scanning inside compressed files and data encoding schemes. If you have enabled Decryption on the firewall, the profile also enables scanning of decrypted content.

The default profile inspects all of the listed protocol decoders for viruses, and generates alerts for SMTP, IMAP, and POP3 protocols while blocking for FTP, HTTP, and SMB protocols. Customized profiles can be used to minimize antivirus inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms.

The Palo Alto Networks WildFire system also provides signatures for persistent threats that are more evasive and have not yet been discovered by other antivirus solutions. As threats are discovered by WildFire, signatures are quickly created and then integrated into the standard antivirus signatures that can be downloaded by Threat Prevention subscribers on a daily basis (sub-hourly for WildFire subscribers).

 © Palo Alto Networks, Inc.

## Anti-spyware Profiles

Anti-spyware profiles blocks spyware on compromised hosts from trying to phone-home or beacon out to external command-and-control (C2) servers, allowing you to detect malicious traffic leaving the network from infected clients. You can apply various levels of protection between zones. For example, you may want to have custom anti-spyware profiles that minimize inspection between trusted zones, while maximizing inspection on traffic received from an untrusted zone, such as Internet facing zones.

You can define your own custom Anti-spyware profiles, or choose one of the following predefined profiles when applying anti-spyware to a security policy:

- **Default**—Uses the default action for every signature, as specified by Palo Alto Networks when the signature is created.
- **Strict**—Overrides the default action of critical, high, and medium severity threats to the block action, regardless of the action defined in the signature file. This profile still uses the default action for medium and informational severity signatures.

In addition, you can enable the DNS Sinkholing action in anti-spyware profiles to enable the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. This feature helps to identify infected hosts on the protected network using DNS traffic Infected hosts can then be easily identified in the traffic and threat logs because any host that attempts to connect to the sinkhole IP address are most likely infected with malware.

Anti-spyware and vulnerability protection profiles are configured similarly.

## Vulnerability Protection Profiles

Vulnerability protection profiles stop attempts to exploit system flaws or gain unauthorized access to systems. While anti-spyware profiles help identify infected hosts as traffic leaves the network, vulnerability protection profiles protect against threats entering the network. For example, vulnerability protection profiles help protect against buffer overflows, illegal code execution, and other attempts to exploit system vulnerabilities. The default vulnerability protection profile protects clients and servers from all known critical, high, and medium-severity threats. You can also create exceptions, which allow you to change the response to a specific signature.

 © Palo Alto Networks, Inc.

## URL Filtering Profiles

URL Filtering profiles enable you to monitor and control how users access the web over HTTP and HTTPS. The firewall comes with a default profile that is configured to block websites such as known malware sites, phishing sites, and adult content sites. You can use the default profile in a security policy, clone it to be used as a starting point for new URL filtering profiles, or add a new URL profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories.

# Data Filtering Profiles

Data filtering profiles prevent sensitive information such as credit card or social security numbers from leaving a protected network. The data filtering profile also allows you to filter on key words, such as a sensitive project name or the word confidential. It is important to focus your profile on the desired file types to reduce false positives. For example, you may only want to search Word documents or Excel spreadsheets. You may also only want to scan web-browsing traffic, or FTP.

You can use default profiles, or create custom data patterns. There are two default profiles:

- CC# (Credit Card)—Identifies credit card numbers using a hash algorithm. The content must match the hash algorithm in order for data to be detected as a credit card number. This method will reduce false positives.
- SSN# (Social Security Number)—Uses an algorithm to detect nine digit numbers, regardless of format. There are two fields: SSN# and SSN# (no dash).

## Weight and Threshold Values

It is important to understand how the weight of an object (SSN, CC#, pattern) is calculated in order to set the appropriate threshold for a condition you are trying to filter. Each occurrence multiplied by the weight value will be added together in order to reach an action threshold (alert or block).

### Example: Filter for Social Security Numbers Only

For simplicity, if you only want to filter files with Social Security Numbers (SSN) and you define a weight of 3 for SSN#, you would use the following formula: each instance of a SSN x weight = threshold increment. In this case, if a Word document has 10 social security numbers you multiply that by the weight of 3, so 10 x 3 = 30. In order to take action for a file that contains 10 social security numbers you would set the threshold to 30. You may want to set an alert at 30 and then block at 60. You may also want to set a weight in the field SSN# (no dash) for Social Security Numbers that do not contain dashes. If multiple settings are used, they will accumulate to reach a given threshold.

### Example: Filter for Social Security Numbers and a Custom Pattern

In this example, we will filter on files that contain Social Security Numbers and the custom pattern confidential. In other words, if a file has Social Security Numbers in addition to the word confidential and the combined instances of those items hit the threshold, the file will trigger an alert or block, depending on the action setting.

SSN# weight = 3

Custom Pattern confidential weight = 20

The custom pattern is case sensitive.

 © Palo Alto Networks, Inc.

If the file contains 20 Social Security Numbers and a weight of 3 is configured, that is 20 x 3 = 60. If the file also contains one instance of the term confidential and a weight of 20 is configured, that is 1 x 20 = 20 for a total of 80. If your threshold for block is set to 80, this scenario would block the file. The alert or block action will be triggered as soon as the threshold is hit.

© Palo Alto Networks, Inc.

## File Blocking Profiles

The firewall uses file blocking profiles two ways: to forward files to WildFire for analysis or to block specified file types over specified applications and in the specified session flow direction (inbound/outbound/both). You can set the profile to alert or block on upload and/or download and you can specify which applications will be subject to the file blocking profile. You can also configure custom block pages that will appear when a user attempts to download the specified file type. This allows the user to take a moment to consider whether or not they want to download a file.

You can configure a file blocking profile with the following actions:

- **Alert**—When the specified file type is detected, a log is generated in the data filtering log.
- **Block**—When the specified file type is detected, the file is blocked and a customizable block page is presented to the user. A log is also generated in the data filtering log.
- **Continue**—When the specified file type is detected, a customizable response page is presented to the user. The user can click through the page to download the file. A log is also generated in the data filtering log. Because this type of forwarding action requires user interaction, it is only applicable for web traffic.
- **Forward**—When the specified file type is detected, the file is sent to WildFire for analysis. A log is also generated in the data filtering log.
- **Continue-and-forward**—When the specified file type is detected, a customizable continuation page is presented to the user. The user can click through the page to download the file. If the user clicks through the continue page to download the file, the file is sent to WildFire for analysis. A log is also generated in the data filtering log.

 © Palo Alto Networks, Inc.

## DoS Protection Profile

DoS protection profiles provide detailed control for Denial of Service (DoS) protection policies. DoS policies allow you to control the number of sessions between interfaces, zones, addresses, and countries based on aggregate sessions or source and/or destination IP addresses. There are two DoS protection mechanisms that the Palo Alto Networks firewalls support.

- **Flood Protection**—Detects and prevents attacks where the network is flooded with packets resulting in too many half-open sessions and/or services being unable to respond to each request. In this case the source address of the attack is usually spoofed. See DoS Protection Against Flooding of New Sessions.
- **Resource Protection**— Detects and prevent session exhaustion attacks. In this type of attack, a large number of hosts (bots) are used to establish as many fully established sessions as possible to consume all of a system's resources.

You can enable both types of protection mechanisms in a single DoS protection profile.

The DoS profile is used to specify the type of action to take and details on matching criteria for the DoS policy. The DoS profile defines settings for SYN, UDP, and ICMP floods, can enable resource protect and defines the maximum number of concurrent connections. After you configure the DoS protection profile, you then attach it to a DoS policy.

When configuring DoS protection, it is important to analyze your environment in order to set the correct thresholds and due to some of the complexities of defining DoS protection policies, this guide will not go into detailed examples. For more information, refer to the Threat Prevention Tech Note.

## Zone Protection Profiles

Zone protection profiles provide additional protection between specific network zones in order to protect the zones against attack. The profile must be applied to the entire zone, so it is important to carefully test the profiles in order to prevent issues that may arise with the normal traffic traversing the zones. When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session. For more information, refer to the Threat Prevention Tech Note.

 © Palo Alto Networks, Inc.

# Security Profile Group

A security profile group is a set of security profiles that can be treated as a unit and then easily added to security policies. Profiles that are often assigned together can be added to profile groups to simplify the creation of security policies. You can also setup a default security profile group—new security policies will use the settings defined in the default profile group to check and control traffic that matches the security policy. Name a security profile group *default* to allow the profiles in that group to be added to new security policies by default. This allows you to consistently include your organization's preferred profile settings in new policies automatically, without administrators having to manually add them each time they create a new policy.

The following sections show how to create a security profile group and how to enable a profile group to be used by default in new security policies:

- Create a Security Profile Group
- Set Up or Override a Default Security Profile Group

## Create a Security Profile Group

Use the following steps to create a security profile group and add it to a security policy.

| **Create a Security Profile Group** | |
|---|---|
| Step 1 Create a security profile group. | 1. Select **Objects > Security Profile Groups** and **Add** a new security profile group.<br>2. Give the profile group a descriptive **Name**, for example, Threats.<br>3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.<br>4. Add existing profiles to the group.<br> 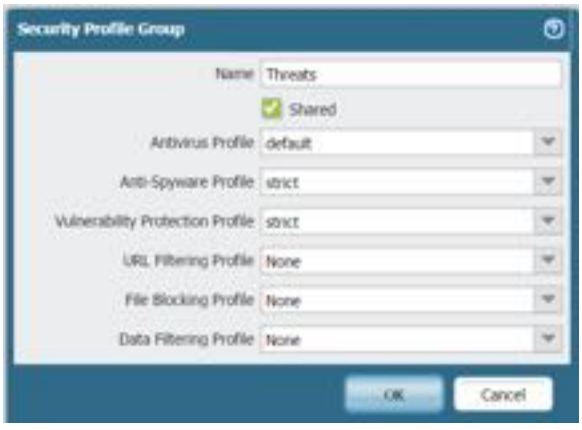 <br>5. Click **OK** to save the profile group. |

| Create a Security Profile Group | | |
|---|---|---|
| Step 2 | Add a security profile group to a security policy. | 1. Select **Policies > Security** and **Add** or modify a security policy rule.<br>2. Select the **Actions** tab.<br>3. In the Profile Setting section, select **Group** for the **Profile Type**.<br>4. In the **Group Profile** drop-down, select the group you created (for example, select the Threats group):<br> 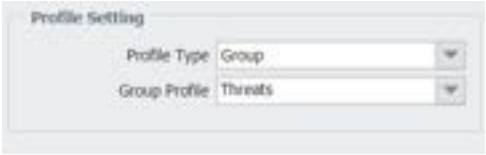 <br>5. Click **OK** to save the policy and **Commit** your changes. |
| Step 3 | Save your changes. | Click **Commit**. |

## Set Up or Override a Default Security Profile Group

Use the following options to set up a default security profile group to be used in new security policies, or to override an existing default group. When an administrator creates a new security policy, the default profile group will be automatically selected as the policy's profile settings, and traffic matching the policy will be checked according to the settings defined in the profile group (the administrator can choose to manually select different profile settings if desired). Use the following options to set up a default security profile group or to override your default settings.

If no default security profile exists, the profile settings for a new security policy are set to **None** by default.

 © Palo Alto Networks, Inc.

| Set Up or Override a Default Security Profile Group | |
|---|---|
| • Create a security profile group. | 1. Select **Objects > Security Profile Groups** and Add a new security profile group.<br>2. Give the profile group a descriptive **Name**, for example, Threats.<br>3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.<br>4. Add existing profiles to the group. For details on creating profiles, see Security Profiles.<br> 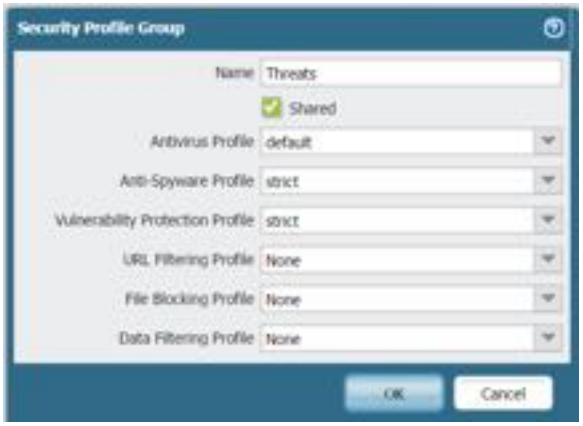 <br>5. Click **OK** to save the profile group.<br>6. **Add the security profile group to a security policy.**<br>7. **Add** or modify a security policy rule and select the **Actions** tab.<br>8. Select **Group** for the **Profile Type**.<br>9. In the **Group Profile** drop-down, select the group you created (for example, select the Threats group):<br> 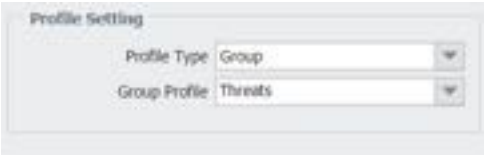 <br>10. Click **OK** to save the policy and **Commit** your changes. |

| Set Up or Override a Default Security Profile Group | |
|---|---|
| • Set up a default security profile group. | 1. Select **Objects > Security Profile Groups** and add a new security profile group or modify an existing security profile group.<br>2. **Name** the security profile group *default*:<br> 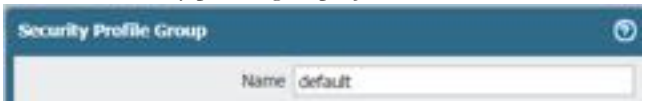 <br>3. Click **OK** and **Commit**.<br>4. Confirm that the *default* security profile group is included in new security policies by default:<br>a. Select **Policies > Security** and **Add** a new security policy.<br>b. Select the **Actions** tab and view the **Profile Setting** fields:<br> <br>By default, the new security policy correctly shows the **Profile Type** set to Group and the *default* **Group Profile** is selected. |
| • Override a default security profile group. | If you have an existing default security profile group, and you do not want that set of profiles to be attached to a new security policy, you can continue to modify the Profile Setting fields according to your preference. Begin by selecting a different Profile Type for your policy (**Policies > Security > Security Policy Rule > Actions**). |

 © Palo Alto Networks, Inc.

# Enumeration of Rules Within a Rulebase

Each rule within a rulebase is automatically numbered and the ordering adjusts as rules are moved or reordered. When filtering rules to find rules that match the specified filter(s), each rule is listed with its number in the context of the complete set of rules in the rulebase and its place in the evaluation order.

On Panorama, pre-rules and post-rules are independently numbered. When rules are pushed from Panorama to a managed firewall, the rule numbering incorporates hierarchy in pre-rules, device rules, and post-rules within a rulebase and reflects the rule sequence and its evaluation order. The **Preview** option in Panorama offers an ordered list view of the total number of rules on a managed device.

**View the Ordered List of Rules Within a Rulebase**

- View the numbered list of rules on the firewall.

  Select **Policies** and any rulebase under it. For example, **Policies > QoS**. The left-most column in the table displays the rule number.

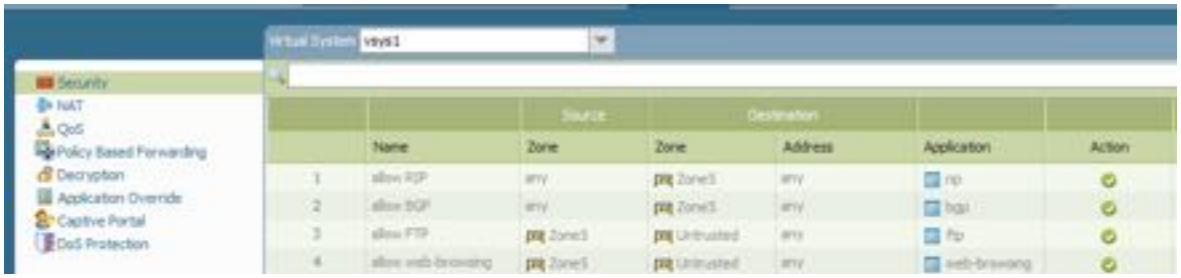


- View the numbered list of rules on Panorama.

  Select **Policies** and any rulebase under it. For example, **Policies > Security> Pre-rules**.

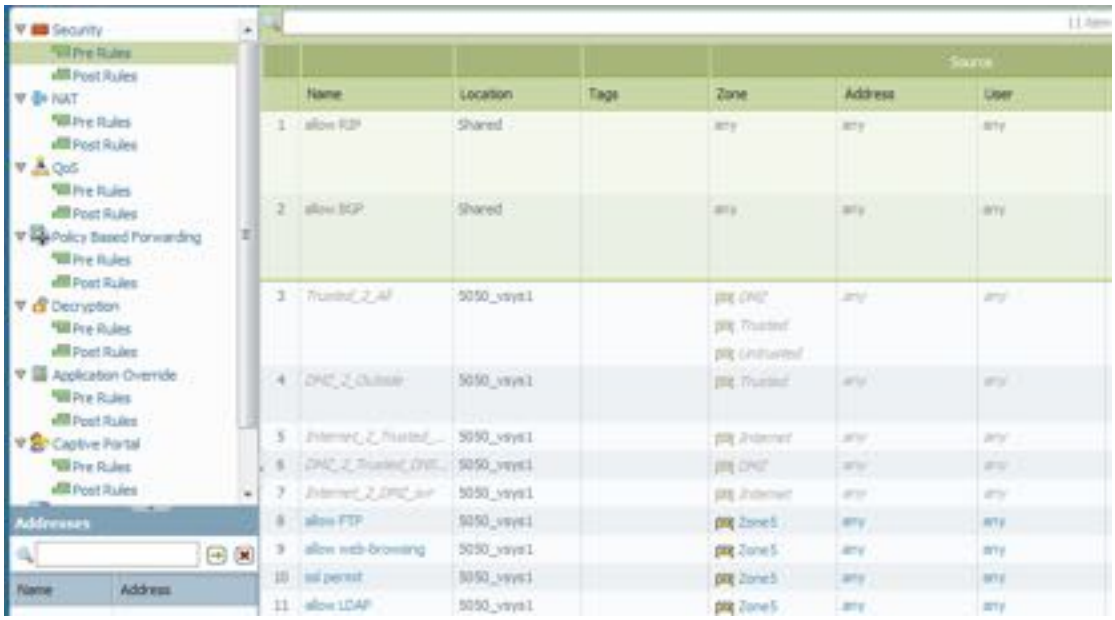

**View the Ordered List of Rules Within a Rulebase (Continued)**

- After you push the rules from Panorama, view the complete list of rules with numbers on the managed device.

  From the web interface of the managed device, select **Policies** and pick any rulebase under it. For example, select **Policies > Security** and view the complete set of numbered rules that will be evaluated on the device.

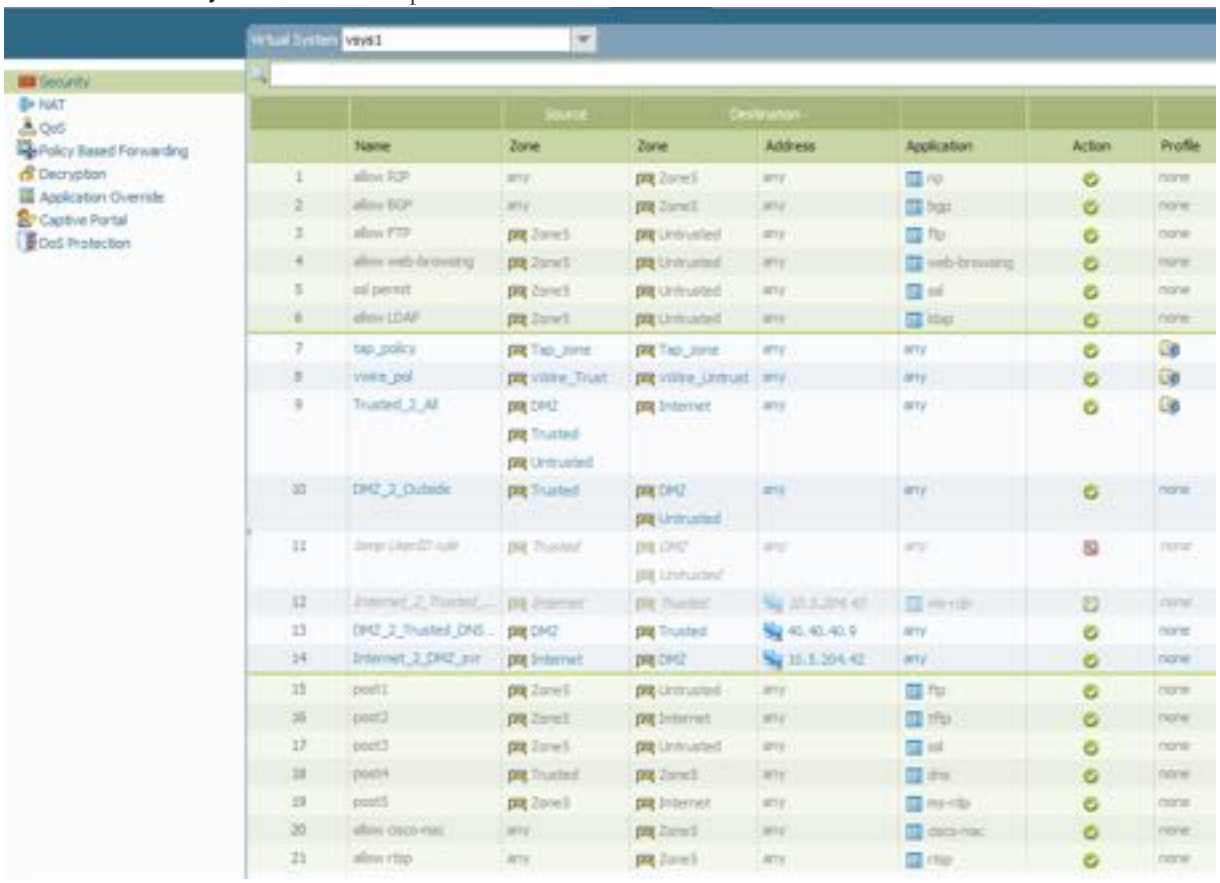


 © Palo Alto Networks, Inc.

# Use Tags to Visually Distinguish Objects

You can tag objects and add color to the tag in order to visually distinguish tagged objects. Tags can be added to the following objects: address objects, address groups, zones, service groups, and policy rules.

The firewall and Panorama support both static tags and dynamic tags, dynamic tags are registered from a variety of sources and are not displayed with the static tags, because dynamic tags are not part of the device configuration. See Register IP Addresses and Tags Dynamically for information on registering tags dynamically. The tags discussed in this section are statically added and are part of the device configuration.

One or more tags can be applied to objects and to policy rules; a maximum of 64 tags can be applied to an object. Panorama supports a maximum of 10,000 tags that can be apportioned across Panorama (shared and device groups) and the managed devices (including devices with multiple virtual systems).

| **Tag Objects** | |
|---|---|
| Step 1 Create tags.<br><br>To tag a zone, you must create a tag with the same name as the zone. When the zone is attached in policy rules, the tag color automatically displays as the background color against the zone name. | 1. Select **Objects > Tags**.<br>2. On Panorama or a multiple virtual system firewall, select the **Device Group** or the **Virtual System** to which this object must belong.<br>3. Click **Add** and enter a **Name** to identify the tag. The maximum length is 127 characters.<br>4. (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.<br>5. (Optional) Assign one of the 16 predefined colors to the tag. By default, no color is selected.<br> <br>6. Click **OK** and **Commit** to save the changes. |
| Step 2 View tags in policy. | 1. Select **Policies** and any rulebase under it.<br>2. Click **Add** to create a policy rule and use the tagged objects you created in Step 1.<br>3. Verify that the tags are in use.<br> |

| Tag Objects (Continued) | |
|---|---|
| Step 3 Working with tags. | • Select **Objects > Tags** to perform any of the following operations with tags:<br>  • Click the link in the **Name** column to edit the properties of a tag.<br>  • Select a tag in the table, and click **Delete** to remove the tag from the firewall.<br>  • Click **Clone** to create a duplicate tag with the same properties. A numerical suffix is added to the tag name. For example, FTP-1.<br>• To apply a tag to an address object, address group, service or service group:<br>  • Create the object. Click **Add**. For example to create a service group, select **Objects > Service Groups > Add**.<br>  • Select the tag(s) from the **Tag** drop-down or enter a phrase to create a new tag.<br> To edit a tag or add color to the tag, select **Objects >Tags**. |

 © Palo Alto Networks, Inc.

# Tips and Tricks for Searching for Objects in Security Policy

The policy filtering mechanism is a powerful tool for searching for objects in even the most complex security policy rulebase. The following topics show examples to help you learn to search the security policy rulebase (**Policies > Security**) efficiently so that you can streamline your management tasks:


- Search Security Policy Rules for Security Profiles
- Search for Disabled Rules in the Security Policy
- Search for Log Forwarding Details in Security Policy
- Search for Log Details in Security Policy
- Search for Scheduled Rules in Security Policy


## Search Security Policy Rules for Security Profiles


- Search for an Antivirus Profile
- Search for a Specific URL Filtering Profile
- Search for an Anti-Spyware Profile
- Search for a Vulnerability Protection Profile
- Search for a File Blocking Profile
- Search for a Data Filtering Profile


### Search for an Antivirus Profile

Use the `profile-setting/profiles/virus/member` syntax to search for an Antivirus security profile. For example, suppose you have an Antivirus profile named **BlockAll** that you use to block viruses in all FTP, HTTP, IMAP, POP3, SMB, and SMTP application traffic. To search for all security policy rules that include this Antivirus profile you would enter the following search string in the filter text box on the **Policy** tab:

```
profile-setting/profiles/virus/member eq BlockAll
```

The search filter is case sensitive. Therefore, when filtering you must enter the object name exactly as it displays on the **Objects** tab. In this case, filtering on `blockall` or `Blockall` would not yield a match because the exact capitalization of the object is `BlockAll`.

After applying the search filter, the rulebase now only displays security policy rules with the BlockAll Antivirus security profile attached as follows:

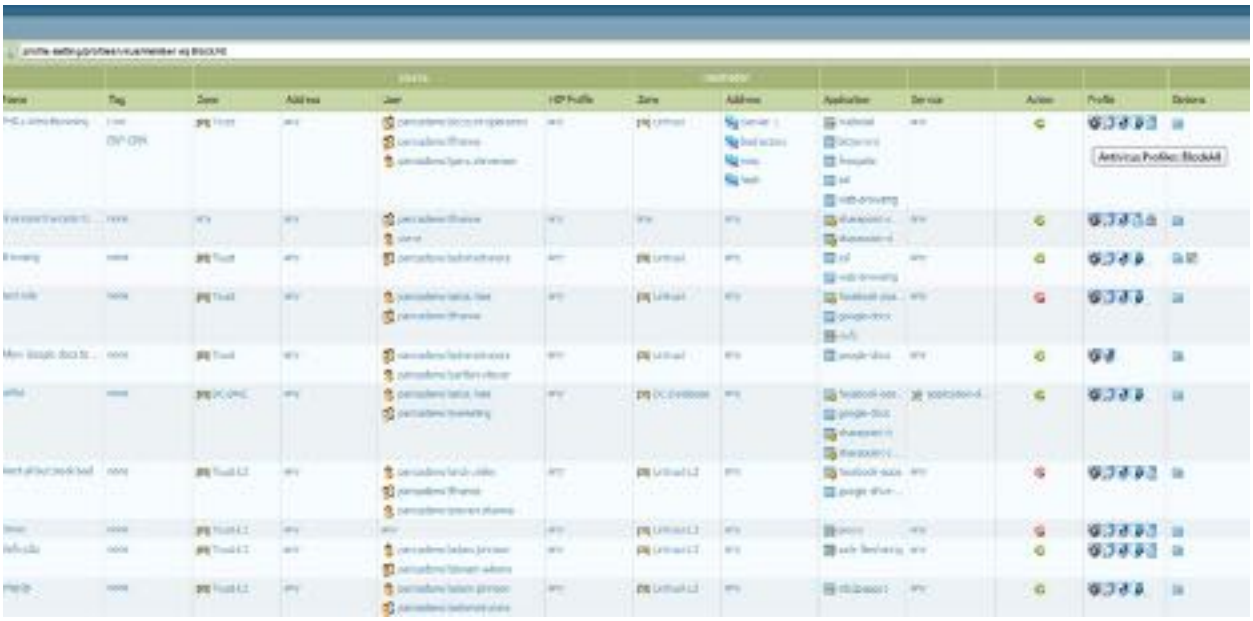

## Search for a Specific URL Filtering Profile

Use the `profile-setting/profiles/url-filtering/member` syntax to search for an antivirus security profile. For example, to search for all security policy rules that have a URL Filtering profile named **block malware**, you would enter the following search string in the filter text box:

```
profile-setting/profiles/url-filtering/member eq "block malware"
```

In this case, because the profile name contains a space, you must enclose the name in quotation marks in order for the filter to match policy rules with the block malware URL Filtering profile attached.

After applying the search filter, the rulebase now only displays security policy rules with the block malware URL Filtering security profile attached as follows:

 © Palo Alto Networks, Inc.

## Search for an Anti-Spyware Profile

Use the `profile-setting/profiles/spyware/member` syntax to search for an Anti-Spyware security profile. For example, to search for all security policy rules that have the **default** Anti-Spyware security profile attached, you would enter the following search string in the filter text box:

```
profile-setting/profiles/spyware/member eq default
```

After applying the search filter, the rulebase now only displays security policy rules with the default Anti-Spyware security profile attached as follows:

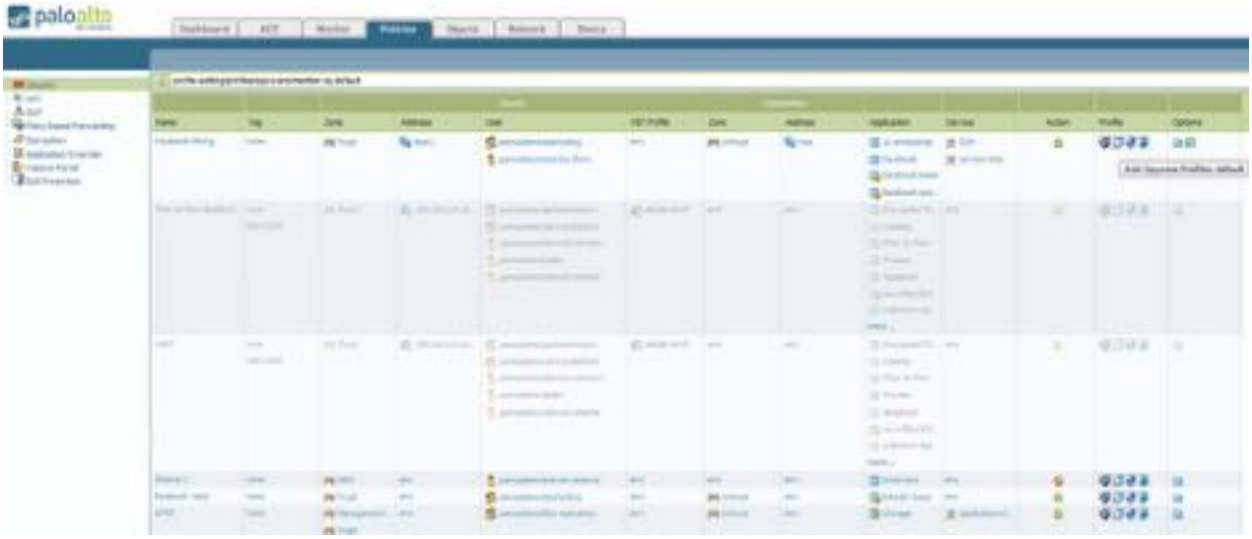

## Search for a Vulnerability Protection Profile

Use the `profile-setting/profiles/vulnerability/member` syntax to search for a Vulnerability Protection security profile. For example, to search for all security policy rules that have the **strict** Vulnerability Protection security profile attached, you would enter the following search string in the filter text box:

```
profile-setting/profiles/vulnerability/member eq strict
```

After applying the search filter, the rulebase now only displays security policy rules with the strict Vulnerability Protection security profile attached as follows:

## Search for a File Blocking Profile

Use the `profile-setting/profiles/file-blocking/member` syntax to search for a File Blocking security profile. For example, suppose you created a File Blocking profile named Wildfire that you will use to forward all PE files to WildFire for analysis. You can quickly verify that all security policy rules have the File Blocking profile attached by searching for all security policy rules that have the **Wildfire** File Blocking security profile attached by entering the following search string in the filter text box:

```
profile-setting/profiles/file-blocking/member eq Wildfire
```

After applying the search filter, the rulebase now displays all security policy rules that have the Wildfire File Blocking security profile attached as follows:

## Search for a Data Filtering Profile

Use the `profile-setting/profiles/data-filtering/member` syntax to search for security policy rules that have a specific Data Filtering security profile attached. For example, suppose you created a Data Filtering profile named **credit cards**. You can search for all security policy rules have this profile attached by searching for all security policy rules that have the **credit cards** Data Filtering security profile attached by entering the following search string in the filter text box:

```
profile-setting/profiles/data-filtering/member eq "credit cards"
```

Notice once again that the profile name must be enclosed in quotation marks because it contains a space.

After applying the search filter, the rulebase now displays only security policy rules that have the credit cards Data Filtering security profile attached as follows:

 © Palo Alto Networks, Inc.

## Search for Disabled Rules in the Security Policy

If you want to get a quick view of all security rules that are currently disabled in your policy, you can display these by entering the following command on the search bar on the Security Policy:

```
disabled eq yes
```

## Search for Log Forwarding Details in Security Policy

If you want to get a quick view of all security rules that are currently sending their logs using a particular log forwarding profile, you can filter the rulebase using the `logsetting` parameter. For example, to filter for all rules that are forwarding their logs using a log forwarding profile named **panorama**, you would enter the following command on the search bar on the **Policies > Security** page:

```
logsetting eq panorama
```

## Search for Log Details in Security Policy

If you want to get a quick view of all security rules that are currently sending logs, you can filter using the log-end and log-start filters. For example, to view all policy rules that generate logs at the end of a session, you can filter the rulebase using the following filter:

```
log-end eq yes
```

Similarly, to view all policy rules that generate logs at the start and the end of a session, you can filter the rulebase as follows:

```
log-start eq yes
```

 © Palo Alto Networks, Inc.

## Search for Scheduled Rules in Security Policy

If you want to get a quick view of all security rules that are set to be active based on a particular time schedule, such as during Lunch time, you can display these by entering the following command on the search bar on the **Policies > Security** page:

```
schedule eq "Lunch time"
```



The search field is case sensitive, so you must enter the object name exactly as it appears in the firewall configuration. In addition, you must enclose object names that contain a space in quotation marks.

# Use a Dynamic Block List in Policy

The firewall or Panorama typically enforce policy for a source or destination IP address that is defined as a static object on the firewall. If you need agility in enforcing policy for a list of source/destination IP addresses that emerge ad hoc, you can use dynamic block lists.

A dynamic block list is a text file that contains a list of IP addresses, IP ranges, or IP subnets, and is hosted on a web server. The dynamic block list can be used to deny or allow access to the IP addresses (IPv4 and IPv6) included in the list. For example, you can use it as a whitelist for allowing a set of IP addresses or as a blacklist to disallow access to the specified IP addresses. At a configured interval, the firewall dynamically imports the list and enforces policy for the IP addresses included in the list. When you modify the list, the firewall retrieves the updates; a configuration change or commit is not required on the firewall. If the web server is unreachable, the firewall or Panorama will use the last successfully retrieved list for enforcing policy until the connection is restored with the web server that hosts the list.


- View the IP Address Limit For Your Firewall Model
- Formatting Guidelines for Dynamic Block Lists
- Enforce Policy with a Dynamic Block List
- View the List of IP addresses in the Dynamic Block List
- Retrieve a Dynamic Block List from Web Server


## View the IP Address Limit For Your Firewall Model

Irrespective of the firewall model, each firewall supports a maximum of 10 Dynamic Block Lists.

To find the maximum number of addresses, address groups, and IP addresses per group, for your model of the firewall, use the following CLI command:

**show system state | match cfg.general.max-address**

For example:

```
admin@PA-7050> show system state | match cfg.general.max-address
cfg.general.max-address: 80000
cfg.general.max-address-group: 8000
cfg.general.max-address-per-group: 500
```

Each list can contain the maximum number of addresses supported by your firewall model minus 300. Up to 300 IP addresses are reserved for internal use on the firewall and are deducted from the available limit. Therefore, in the example above, the firewall can have a maximum of 79,700 IP addresses.

 © Palo Alto Networks, Inc.

## Formatting Guidelines for Dynamic Block Lists

The dynamic block list can include individual IP addresses, subnet addresses (address/mask), or range of IP addresses. In addition, the block list can include comments and special characters such as `*` , `:` , `;` , `#`, or `/`. The syntax for each line in the list is `[IP address, IP/Mask, or IP start range-IP end range] [space] [comment]`.

Because the firewall ignores incorrectly formatted lines, use these guidelines when defining the list:

- Enter each IP address/range/subnet in a new line; URLs are not supported in this list.
- If you add comments, the comment must be on the same line as the IP address/range/subnet. The space at the end of the IP address is the delimiter that separates a comment from the IP address.

An example:

```
192.168.20.10/32
2001:db8:123:1::1 #test IPv6 address
192.168.20.0/24 ; test internal subnet
2001:db8:123:1::/64 test internal IPv6 range
192.168.20.40-192.168.20.50
```

For an IP address that is blocked, you can display a notification page only if the protocol is HTTP.

## Enforce Policy with a Dynamic Block List

| **Enforce Policy with a Dynamic Block List** | | |
|---|---|---|
| Step 1 | Create the dynamic block list and host it on a web server, so that the firewall can retrieve the list for policy evaluation. | 1. Create a text file and enter the IP addresses for which you want to enforce policy. For syntax, see Formatting Guidelines for Dynamic Block Lists. |
| Step 2 | Create a dynamic block list object on the firewall. | 1. Select **Objects > Dynamic Block Lists**.<br>2. Click **Add** and enter a descriptive **Name** for the list.<br>3. (Optional) Select **Shared**, to share the list with all virtual systems on a device that is enabled for multiple virtual systems. By default, the object is created on the virtual system that is currently selected in the **Virtual Systems** drop-down.<br>4. Enter the **Source** URL (hostname or IP address and the path) for the list you just created on the web server. For example, https://1.2.3.4/DBL_2014<br>5. Click **Test Source URL** to verify that the firewall or Panorama can connect to the web server.<br>6. (Optional) Specify the **Repeat** frequency at which the firewall or Panorama must retrieve the list. By default the list is retrieved ever hour.<br>7. Click **OK** to save the changes. |

| Enforce Policy with a Dynamic Block List | | |
|---|---|---|
| Step 3 | Use the dynamic block list as a source or destination address object in policy.<br> Create separate dynamic block lists if you want to specify allow and deny actions for specific IP addresses.<br>The list can be referenced in any policy type. In this example, we attach it as a destination object in security policy. | 1. Select **Policies > Security**.<br>2. Click **Add** and give the rule a descriptive name in the **General** tab.<br>3. In the **Source** tab, select the **Source Zone**.<br>4. In the **Destination** tab, select the **Destination Zone** and select the dynamic block list as the Destination Address.<br>5. In the **Service/ URL Category** tab, make sure the **Service** is set to **application-default**.<br>6. In the **Actions** tab, set the **Action Setting** to Allow or Deny.<br>7. Leave all the other options at the default values.<br>8. Click **OK** to save the changes.<br>9. **Commit** the changes. |
| Step 4 | Test that the policy action is enforced. | 1. Access a IP address that is included in the dynamic block list and verify that action you defined is enforced.<br>2. Select **Monitor > Logs > Traffic** and see the log entry for the session.<br>3. To verify the policy rule that matches a flow, use the following CLI command:<br>**test security-policy-match source** *<IP_address>* **destination** *<IP_address>* **destination port** *<port_number>* **protocol** *<protocol_number>* |

## View the List of IP addresses in the Dynamic Block List

**View the IP Addresses Included in the Dynamic Block List**

To view the list of IP addresses that the firewall has retrieved from the web server enter the following CLI command:

**request system external-list show name *<name>***

For example, for a list named case DBL_2014, the output is:

```
vsys1/DBL_2014:
  Next update at: Wed Aug 27 16:00:00 2014
  IPs:
  1.1.1.1
  1.2.2.2/20 #test China
  192.168.255.0; test internal
  192.168.254.0/24 test internal range
```

 © Palo Alto Networks, Inc.

## Retrieve a Dynamic Block List from Web Server

The firewall or Panorama can be configured to retrieve the list from the web server on an hourly, daily, weekly, or monthly basis. If you have added or deleted IP addresses on the list and need to trigger an immediate refresh, you must use the Command Line Interface.

**Retrieve a Dynamic Block List**

1. Enter the command: **request system external-list refresh name *<name>***

   For example, `request system external-list refresh name DBL_2014`
2. Get the job ID for the refresh job using the CLI command: **show jobs all**

   Look for the last EBL Refresh job in the list.
3. View the details for the job ID. Use the command **show jobs id *<number>***

   A message indicating the success or failure displays. For example:

```
admin@PA-200> show jobs id 55
Enqueued                    ID             Type    Status Result Completed
------------------------------------------------------------------------------
2014/08/26 15:34:14         55       EBLRefresh       FIN     OK 15:34:40
Warnings:
Details:
```

# Register IP Addresses and Tags Dynamically

To mitigate the challenges of scale, lack of flexibility and performance, the architecture in networks today allows for clients, servers and applications to be provisioned, changed, and deleted on demand. This agility poses a challenge for security administrators because they have limited visibility into the IP addresses of the dynamically provisioned clients, servers and the plethora of applications that can be enabled on these virtual resources.

The firewall (hardware-based platforms and the VM-Series) supports the ability to register IP addresses and tags dynamically. The IP addresses and tags can be registered on the firewall directly or registered on the firewall through Panorama.This dynamic registration process can be enabled using any of the following options:

- **User-ID agent for Windows**—In an environment where you've deployed the User-ID agent, you can enable the User-ID agent to monitor up to 100 VMware ESXi and/or vCenter Servers. As you provision or modify virtual machines on these VMware servers, the agent can retrieve the IP address changes and share them with the firewall.
- **VM Information Sources**—Allows you to monitor VMware ESXi and vCenter Server, and the AWS-VPC to retrieve IP address changes when you provision or modify virtual machines on these sources. VM Information Sources polls for a predefined set of attributes and does not require external scripts to register the IP addresses through the XML API. See Monitor Changes in the Virtual Environment.
- **VMware Service Manager** (only available for the integrated NSX solution)—The integrated NSX solution is designed for automated provisioning and distribution of Palo Alto Networks next-generation security services and the delivery of dynamic context-based security policies using Panorama. The NSX Manager updates Panorama with the latest information on the IP addresses and tags associated with the virtual machines deployed in this integrated solution. For information on this solution, see Set Up a VM-Series NSX Edition Firewall.
- **XML API**—The firewall and Panorama support an XML API that uses standard HTTP requests to send and receive data. You can use this API to register IP addresses and tags with the firewall or Panorama. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports REST-based services. Refer to the PAN-OS XML API Usage Guide for details.

For information on creating and using Dynamic Address Groups, see Use Dynamic Address Groups in Policy.

For the CLI commands for registering tags dynamically, see CLI Commands for Dynamic IP Addresses and Tags.

 © Palo Alto Networks, Inc.

# Monitor Changes in the Virtual Environment

To secure applications and prevent threats in an environment where new users and servers are constantly emerging, your security policy must be nimble. To be nimble, the firewall must be able to learn about new or modified IP addresses and consistently apply policy without requiring configuration changes on the firewall.

This capability is provided by the coordination between the **VM Information Sources** and **Dynamic Address Groups** features on the firewall. The firewall and Panorama provide an automated way to gather information on the virtual machine (or guest) inventory on each monitored source and create policy objects that stay in sync with the dynamic changes on the network.


- Enable VM Monitoring to Track Changes on the Virtual Network
- Attributes Monitored in the AWS and VMware Environments
- Use Dynamic Address Groups in Policy

## Enable VM Monitoring to Track Changes on the Virtual Network

VM information sources provides an automated way to gather information on the Virtual Machine (VM) inventory on each monitored source (host); the firewall can monitor the VMware ESXi and vCenter Server, and the AWS-VPC. As virtual machines (guests) are deployed or moved, the firewall collects a predefined set of attributes (or metadata elements) as tags; these tags can then be used to define Dynamic Address Groups (see Use Dynamic Address Groups in Policy) and matched against in policy.

Up to 10 VM information sources can be configured on the firewall or pushed using Panorama templates. By default, the traffic between the firewall and the monitored sources uses the management (MGT) port on the firewall.

**VM Information Sources** offers easy configuration and enables you to monitor a predefined set of 16 metadata elements or attributes. See Attributes Monitored in the AWS and VMware Environmentsfor the list.

 © Palo Alto Networks, Inc.

**Set up the VM Monitoring Agent**

Step 1 Enable the VM Monitoring Agent.

Up to 10 sources can be configured for each firewall, or for each virtual system on a multiple virtual systems capable firewall.

If your firewalls are configured in a high availability configuration:

- An active/passive setup, only the active firewall monitors the VM sources.
- An active/active setup, only the firewall with the priority value of primary monitors the VM sources.

1. Select **Device > VM Information Sources**.
2. Click **Add** and enter the following information:
   - A **Name** to identify the VMware ESX(i) or vCenter Server that you want to monitor.
   - Enter the **Host information for the server—**hostname or IP address and the **Port** on which it is listening.
   - Select the **Type** to indicate whether the source is a **VMware ESX(i)** server or a **VMware vCenter** Server.
   - Add the credentials (**Username** and **Password**) to authenticate to the server specified above.
   - Use the credentials of an administrative user to enable access.
   - (Optional) Modify the **Update interval** to a value between 5-600 seconds. By default, the firewall polls every 5 seconds. The API calls are queued and retrieved within every 60 seconds, so updates may take up to 60 seconds plus the configured polling interval.

   - (Optional) Enter the interval in hours when the connection to the monitored source is closed, if the host does not respond. (default: 2 hours, range 2-10 hours)
     To change the default value, select the check box to **Enable timeout when the source is disconnected** and specify the value. When the specified limit is reached or if the host cannot be accessed or does not respond, the firewall will close the connection to the source.
   - Click **OK**, and **Commit** the changes.
   - Verify that the connection **Status** displays as connected .

| Set up the VM Monitoring Agent (Continued) | |
|---|---|
| Step 2 Verify the connection status. | Verify that the connection **Status** displays as connected.<br> <br>If the connection status is pending or disconnected, verify that the source is operational and that the firewall is able to access the source. If you use a port other than the MGT port for communicating with the monitored source, you must change the service route (**Device > Setup > Services**, click the **Service Route Configuration** link and modify the **Source Interface** for the **VM Monitor** service). |

 © Palo Alto Networks, Inc.

## Attributes Monitored in the AWS and VMware Environments

Each VM on a monitored ESXi or vCenter server must have VMware Tools installed and running. VMware Tools provide the capability to glean the IP address(es) and other values assigned to each VM.

In order to collect the values assigned to the monitored VMs, the firewall monitors the following predefined set of attributes:

| Attributes Monitored on a VMware Source | Attributes Monitored on the AWS-VPC |
|---|---|
| • UUID | • Architecture |
| • Name | • Guest OS |
| • Guest OS | • Image ID |
| • VM State — the power state can be poweredOff, poweredOn, standBy, and unknown. | • Instance ID |
| • Annotation | • Instance State |
| • Version | • Instance Type |
| • Network — Virtual Switch Name, Port Group Name, and VLAN ID | • Key Name |
| • Container Name —vCenter Name, Data Center Object Name, Resource Pool Name, Cluster Name, Host, Host IP address. | • Placement—Tenancy, Group Name, Availability Zone |
| | • Private DNS Name |
| | • Public DNS Name |
| | • Subnet ID |
| | • Tag (key, value) (up to5 tags supported per instance |
| | • VPC ID |

## Use Dynamic Address Groups in Policy

Dynamic address groups are used in policy. They allow you to create policy that automatically adapts to changes—adds, moves, or deletions of servers. It also enables the flexibility to apply different rules to the same server based on *tags* that define its role on the network, the operating system, or the different kinds of traffic it processes.

A dynamic address group uses tags as a filtering criteria to determine its members. The filter uses logical *and* and *or* operators. All IP addresses or address groups that match the filtering criteria become members of the dynamic address group. Tags can be defined statically on the firewall and/or registered (dynamically) to the firewall. The difference between static and dynamic tags is that static tags are part of the configuration on the firewall, and dynamic tags are part of the runtime configuration. This implies that a commit is not required to update dynamic tags; the tags must however be used by Dynamic Address Groups that are referenced in policy, and the policy must be committed on the device.

To dynamically register tags, you can use the XML API or the VM Monitoring agent on the firewall or on the User-ID agent. Each tag is a metadata element or attribute-value pair that is registered on the firewall or Panorama. For example, IP1 {tag1, tag2,.....tag32}, where the IP address and the associated tags are maintained as a list; each registered IP address can have up to 32 tags such as the operating system, the datacenter or the virtual switch to which it belongs. Within 60 seconds of the API call, the firewall registers the IP address and associated tags, and automatically updates the membership information for the dynamic address group(s).

The maximum number of IP addresses that can be registered for each platform is different. Use the following table for specifics on your platform:

| Platform | Maximum number of dynamically registered IP addresses |
|---|---|
| PA-7050, PA-5060, VM-1000-HV | 100,000 |
| PA-5050 | 50,000 |
| PA-5020 | 25,000 |
| PA-4000 Series, PA-3000 Series | 5000 |
| PA-2000 Series, PA-500, PA-200, VM-300, VM-200, VM-100 | 1000 |

The following example shows how dynamic address groups can simplify network security enforcement. The example workflow shows how to:

- Enable the VM Monitoring agent on the firewall, to monitor the VMware ESX(i) host or vCenter Server and register VM IP addresses and the associated tags.
- Create dynamic address groups and define the tags to filter. In this example, two address groups are created. One that only filters for dynamic tags and another that filters for both static and dynamic tags to populate the members of the group.
- Validate that the members of the dynamic address group are populated on the firewall.
- Use dynamic address groups in policy. This example uses two different security policies:
  - A security policy for all Linux servers that are deployed as FTP servers; this rule matches on dynamically registered tags.

 © Palo Alto Networks, Inc.

– A security policy for all Linux servers that are deployed as web servers; this rule matches on a dynamic address group that uses static and dynamic tags.

- Validate that the members of the dynamic address groups are updated as new FTP or web servers are deployed. This ensure that the security rules are enforced on these new virtual machines too.

| **Use Dynamic Address Groups in Policy** | |
|---|---|
| Step 1 Enable VM Source Monitoring. | See Enable VM Monitoring to Track Changes on the Virtual Network. |
| Step 2 Create dynamic address groups on the firewall.<br><br>View the tutorial to see a big picture view of the feature. | 1. Log in to the web interface of the firewall.<br>2. Select **Object > Address Groups**.<br>3. Click **Add** and enter a **Name** and a **Description** for the address group.<br>4. Select **Type** as **Dynamic**.<br>5. Define the match criteria. You can select dynamic and static tags as the match criteria to populate the members of the group. Click **Add Match Criteria**, and select the **And** or **Or** operator and select the attributes that you would like to filter for or match against. and then click **OK**.<br> <br><br>6. Click **Commit**. |

The match criteria for each dynamic address group in this example is as follows:

ftp_server: matches on the guest operating system “Linux 64-bit” and annotated as “ftp” ('guestos.Ubuntu Linux 64-bit' and 'annotation.ftp').

web-servers: matches on two criteria—the tag black or if the guest operating system is Linux 64-bit and the name of the server us Web_server_Corp. ('guestos.Ubuntu Linux 64-bit' and 'vmname.WebServer_Corp' or 'black')

**Use Dynamic Address Groups in Policy (Continued)**

Step 3 Use dynamic address groups in policy.

View the tutorial.

1. Select **Policies > Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the Dynamic address group you created in Step 2 above.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeats Steps 1 through 6 above to create another policy rule.
8. Click **Commit**.

This example shows how to create two policies: one for all access to FTP servers and the other for access to web servers.

Step 4 Validate that the members of the dynamic address group are populated on the firewall.

1. Select **Policies > Security**, and select the rule.
2. Select the drop-down arrow next to the address group link, and select **Inspect**. You can also verify that the match criteria is accurate.

3. Click the **more** link and verify that the list of registered IP addresses is displayed.

Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.


© Palo Alto Networks, Inc.

# CLI Commands for Dynamic IP Addresses and Tags

The Command Line Interface on the firewall and Panorama give you a detailed view into the different sources from which tags and IP addresses are dynamically registered. It also allows you to audit registered and unregistered tags. The following examples illustrate the capabilities in the CLI.

| Example | CLI Command |
|---|---|
| View all registered IP addresses that match the tag, `state.poweredOn` or that are not tagged as `vSwitch0` | `show log iptag tag_name equal state.poweredOn`<br>`show log iptag tag_name not-equal switch.vSwitch0` |
| View all dynamically registered IP addresses that were sourced by VM Information Source with name `vmware1` and tagged as `poweredOn` | `show vm-monitor source source-name vmware1 tag state.poweredOn registered-ip all`<br><br>`registered IP Tags`<br>`---------- ----------`<br>`fe80::20c:29ff:fe69:2f76 "state.poweredOn"`<br>`10.1.22.100 "state.poweredOn"`<br>`2001:1890:12f2:11:20c:29ff:fe69:2f76 "state.poweredOn"`<br>`fe80::20c:29ff:fe69:2f80 "state.poweredOn"`<br>`192.168.1.102 "state.poweredOn"`<br>`10.1.22.105 "state.poweredOn"`<br>`2001:1890:12f2:11:2cf8:77a9:5435:c0d "state.poweredOn"`<br>`fe80::2cf8:77a9:5435:c0d "state.poweredOn"` |
| Clear all IP addresses and tags learned from a specific VM Monitoring source without disconnecting the source. | `debug vm-monitor clear source-name <name>` |
| Display IP addresses registered from all sources. | `show object registered-ip all` |
| Display the count for IP addresses registered from all sources. | `show object registered-ip all option count` |
| Clear IP addresses registered from all sources | `debug object registered-ip clear all` |
| Add or delete tags for a given IP address that was registered using the XML API. | `debug object test registered-ip [<register/unregister>] <ip/netmask> <tag>` |

| Example | CLI Command |
|---|---|
| View all tags registered from a specific information source. | `show vm-monitor source source-name vmware1 tag all`<br>`vlanId.4095`<br>`vswitch.vSwitch1`<br>`host-ip.10.1.5.22`<br>`portgroup.TOBEUSED`<br>`hostname.panserver22`<br>`portgroup.VM Network 2`<br>`datacenter.ha-datacenter`<br>`vlanId.0`<br>`state.poweredOn`<br>`vswitch.vSwitch0`<br>`vmname.Ubuntu22-100`<br>`vmname.win2k8-22-105`<br>`resource-pool.Resources`<br>`vswitch.vSwitch2`<br>`guestos.Ubuntu Linux 32-bit`<br>`guestos.Microsoft Windows Server 2008 32-bit`<br>`annotation.`<br>`version.vmx-08`<br>`portgroup.VM Network`<br>`vm-info-source.vmware1`<br>`uuid.564d362c-11cd-b27f-271f-c361604dfad7`<br>`uuid.564dd337-677a-eb8d-47db-293bd6692f76`<br>`Total: 22` |
| View all tags registered from a specific data source, for example from the VM Monitoring Agent on the firewall, the XML API, Windows User-ID Agent or the CLI. | • To view tags registered from the CLI:<br>`show log iptag datasource_type equal unknown`<br>• To view tags registered from the XML API:<br>`show log iptag datasource_type equal xml-api`<br>• To view tags registered from VM Information sources:<br>`show log iptag datasource_type equal vm-monitor`<br>• To view tags registered from the Windows User-ID agent:<br>`show log iptag datasource_type equal xml-api datasource_subtype equal user-id-agent` |
| View all tags that are registered for a specific IP address (across all sources). | `debug object registered-ip show tag-source ip` *`ip_address`* `tag all` |

 © Palo Alto Networks, Inc.

# Policy-Based Forwarding

Normally, the firewall uses the destination IP address in a packet to determine the outgoing interface. The firewall uses the routing table associated with the virtual router to which the interface is connected to perform the route lookup. Policy-Based Forwarding (PBF) allows you to override the routing table, and specify the outgoing or *egress* interface based on specific parameters such as source or destination IP address, or type of traffic.


- PBF
- Create a Policy-Based Forwarding Rule
- Use Case: PBF for Outbound Access with Dual ISPs
- Use Case: PBF for Routing Traffic Through Virtual Systems

# PBF

PBF rules allow traffic to take an alternative path from the next hop specified in the route table, and are typically used to specify an egress interface for security or performance reasons. Let's say your company has two links between the corporate office and the branch office: a cheaper Internet link and a more expensive leased line. The leased line is a high-bandwidth, low-latency link. For enhanced security, you can use PBF to send applications that aren't encrypted traffic, such as FTP traffic, over the private leased line and all other traffic over the Internet link. Or, for performance, you can choose to route business-critical applications over the leased line while sending all other traffic, such as web browsing, over the cheaper link.

## Egress Path and Symmetric Return

Using PBF, you can direct traffic to a specific interface on the firewall, drop the traffic, or direct traffic to another virtual system (on systems enabled for multiple virtual systems).

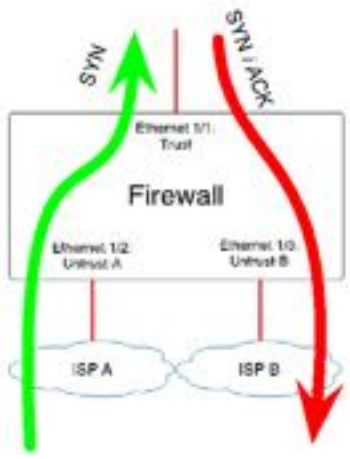


In networks with asymmetric routes, such as in a dual ISP environment, connectivity issues occur when traffic arrives at one interface on the firewall and leaves from another interface. If the route is asymmetrical, where the forward (SYN packet) and return (SYN/ACK) paths are different, the firewall is unable to track the state of the entire session and this causes a connection failure. To ensure that the traffic uses a symmetrical path, which means that the traffic arrives at and leaves from the same interface on which the session was created, you can enable the *Symmetric Return* option.

With symmetric return, the virtual router overrides a routing lookup for return traffic and instead directs the flow back to the MAC address from which it received the SYN packet (or first packet). However, if the destination IP address is on the same subnet as the ingress/egress interface's IP address, a route lookup is performed and symmetric return is not enforced. This behavior prevents traffic from being blackholed.

To determine the next hop for symmetric returns, the firewall uses an Address Resolution Protocol (ARP) table. The maximum number of entries that this ARP table supports is limited by the firewall model and the value is not user configurable. To determine the limit for your model, use the CLI command: **`show pbf return-mac all`**.

## Path Monitoring

Path monitoring allows you to verify connectivity to an IP address so that the firewall can direct traffic through an alternate route, when needed. The firewall uses ICMP pings as *heartbeats* to verify that the specified IP address is reachable.

A monitoring profile allows you to specify the threshold number of heartbeats to determine whether the IP address is reachable. When the monitored IP address is unreachable, you can either disable the PBF rule or specify a *fail-over* or *wait-recover* action. Disabling the PBF rule allows the virtual router to take over the routing decisions.

 © Palo Alto Networks, Inc.

The following table lists the difference in behavior for a path monitoring failure on a new session versus an established session.

| Behavior of a session on a monitoring failure | If the rule stays enabled when the monitored IP address is unreachable<br>☐ Disable this rule if nexthop/monitor ip is unreachable | If rule is disabled when the monitored IP address is unreachable<br>☑ Disable this rule if nexthop/monitor ip is unreachable |
|---|---|---|
| For an established session | **wait-recover**—Continue to use egress interface specified in the PBF rule | **wait-recover**—Continue to use egress interface specified in the PBF rule |
| | **fail-over**—Use path determined by routing table (no PBF) | **fail-over**—Use path determined by routing table (no PBF) |
| For a new session | **wait-recover**—Use path determined by routing table (no PBF) | **wait-recover**—Check the remaining PBF rules. If no match, use the routing table |
| | **fail-over**—Use path determined by routing table (no PBF) | **fail-over**—Check the remaining PBF rules. If no match, use the routing table |

## Service Versus Applications in PBF

PBF rules are applied either on the first packet (SYN) or the first response to the first packet (SYN/ACK). This means that a PBF rule may be applied before the firewall has enough information to determine the application. Therefore, application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application.

However, if you specify an application in a PBF rule, the firewall performs *App-ID caching*. When an application passes through the firewall for the first time, the firewall does not have enough information to identify the application and therefore cannot enforce the PBF rule. As more packets arrive, the firewall determines the application and creates an entry in the App-ID cache and retains this App-ID for the session.When a new session is created with the same destination IP address, destination port, and protocol ID, the firewall could identify the application as the same from the initial session (based on the App-ID cache) and apply the PBF rule. Therefore, a session that is not an exact match and is not the same application, can be forwarded based on the PBF rule.

Further, applications have dependencies and the identity of the application can change as the firewall receives more packets. Because PBF makes a routing decision at the start of a session, the firewall cannot enforce a change in application identity. YouTube, for example, starts as web-browsing but changes to Flash, RTSP, or YouTube based on the different links and videos included on the page. However with PBF, because the firewall identifies the application as web-browsing at the start of the session, the change in application is not recognized thereafter.

PBF rules cannot be based on domain names; only IP addresses are valid; also, you cannot use custom-applications, application-filters or application groups in PBF rules.

# Create a Policy-Based Forwarding Rule

Use a PBF rule to direct traffic to a specific egress interface on the firewall, and override the default path for the traffic.

| Create a PBF Rule | |
|---|---|
| Step 1 Create a PBF rule.<br><br>When creating a PBF rule you must specify a name for the rule, a source zone or interface, and an egress interface. All other components are either optional or have a default value provided. | 1. Select **Policies > Policy Based Forwarding** and click **Add.**<br>2. Give the rule a descriptive name in the **General** tab.<br>3. In the **Source** tab, select the following:<br>a. Select the **Type**—**Zone** or **Interface**— to which the forwarding policy will be applied, and the relevant zone or interface.<br><br>PBF is only supported on Layer 3 interfaces.<br><br>b. (Optional) Specify the **Source Address** to which PBF will apply. For example, a specific IP address or subnet IP address from which you want to forward traffic to the interface or zone specified in this rule.<br><br>Use the **Negate** option to exclude a one or more source IP addresses from the PBF rule. For example, if your PBF rule directs all traffic from the specified zone to the Internet, **Negate** allows you to exclude internal IP addresses from the PBF rule.<br><br>The evaluation order is top down. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated.<br><br>c. (Optional) **Add** and select the **Source User** or groups of users to whom the policy applies.<br>4. In the **Destination/Application/Service** tab, select the following:<br>a. **Destination Address**. By default the rule applies to **Any** IP address. Use the **Negate** option to exclude one or more destination IP addresses from the PBF rule.<br><br>b. Select the Application(s) or Service(s) that you want to control using PBF.<br><br>Application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application. For more details, see Service Versus Applications in PBF. |


© Palo Alto Networks, Inc.

| Create a PBF Rule | |
|---|---|
| | 5. In the **Forwarding** tab, select the following:<br>a. Set the **Action.** The options are as follows:<br>– **Forward**—Directs the packet to a specific **Egress Interface**. Enter the **Next Hop** IP address for the packet.<br>– **Forward To VSYS**—(On a device enabled for multiple virtual systems) Select the virtual system to which to forward the packet.<br>– **Discard**—Drop the packet.<br>– **No PBF**—Exclude the packets that match the criteria for source/destination/application/service defined in the rule. Matching packets use the route table instead of PBF; the firewall uses the route table to exclude the matched traffic from the redirected port.<br>To trigger the specified action at a daily, weekly or non-recurring frequency, create and attach a **Schedule**.<br>b. (Optional) Enable Monitoring to verify connectivity to a target IP address or to the next hop IP address. Select **Monitor** and attach a monitoring **Profile** (default or custom) that specifies the action when the IP address is unreachable.<br>c. (Optional, required for asymmetric routing environments) Select **Enforce Symmetric Return** and enter one or more IP addresses in the **Next Hop Address List**.<br>Enabling symmetric return ensures that return traffic (say, from the Trust zone on the LAN to the Internet) is forwarded out through the same interface through which traffic ingresses from the Internet. |
| Step 2 Save the policies to the running configuration on the device. | Click **Commit**.<br>The PBF rule is in effect. |

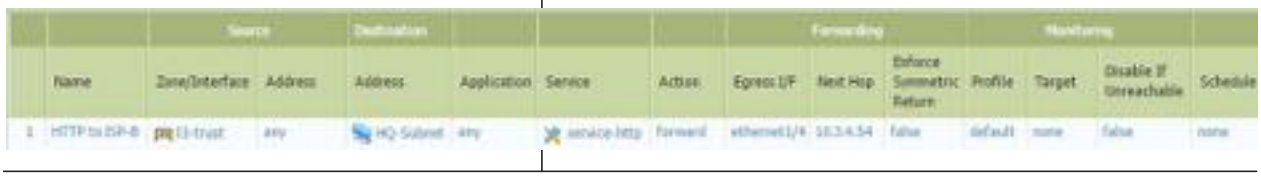

## Use Case: PBF for Outbound Access with Dual ISPs

In this use case, the branch office has a dual ISP configuration and implements PBF for redundant Internet access. The backup ISP is the default route for traffic from the client to the web servers. In order to enable redundant Internet access without using an internetwork protocol such as BGP, we use PBF with destination interface-based source NAT and static routes, and configure the firewall as follows:

- Enable a PBF rule that routes traffic through the primary ISP, and attach a monitoring profile to the rule. The monitoring profile triggers the firewall to use the default route through the backup ISP when the primary ISP is unavailable.
- Define Source NAT rules for both the primary and backup ISP that instruct the firewall to use the source IP address associated with the egress interface for the corresponding ISP. This ensures that the outbound traffic has the correct source IP address.
- Add a static route to the backup ISP, so that when the primary ISP is unavailable, the default route comes into effect and the traffic is directed through the backup ISP.

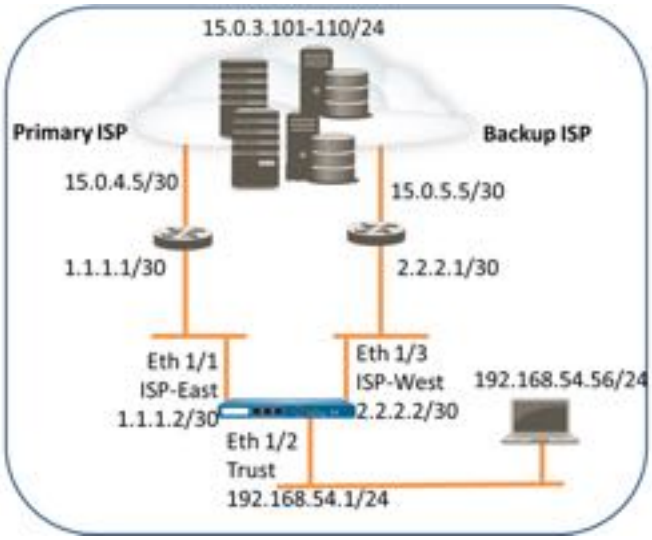


 © Palo Alto Networks, Inc.

| PBF for Outbound Access with Dual ISPs | | |
|---|---|---|
| Step 1 | Configure the ingress and the egress interfaces on the firewall.<br><br>Egress interfaces can be in the same zone. In this example we assign the egress interfaces to different zones. | 1. Select **Network > Interfaces** and then select the interface you want to configure, for example, Ethernet1/1 and Ethernet1/3.<br>The interface configuration on the firewall used in this example is as follows:<br>• Ethernet 1/1 connected to the primary ISP:<br>– Zone: ISP-East<br>– IP Address:1.1.1.2/30<br>– Virtual Router: Default<br>• Ethernet 1/3 connected to the backup ISP:<br>– Zone: ISP-West<br>– IP Address:2.2.2.2/30<br>– Virtual Router: Default<br>• Ethernet 1/2 is the ingress interface, used by the network clients to connect to the Internet:<br>– Zone: Trust<br>– IP Address:192.168. 54.1/24<br>– Virtual Router: Default<br>2. To save the interface configuration, click **OK**. |
| Step 2 | On the virtual router, add a static route to the backup ISP. | 1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog.<br>2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and specify the **Destination** IP address for which you are defining the static route. In this example, we use 0.0.0.0/0 for all traffic.<br>3. Select the **IP Address** radio button and set the **Next Hop** IP address for your router that connects to the backup Internet gateway. In this example, 2.2.2.1.<br>4. Specify a cost metric for the route. In this example, we use 10.<br> <br>5. Click **OK** twice to save the virtual router configuration. |

**PBF for Outbound Access with Dual ISPs**

Step 3 Create a PBF rule that directs traffic to the interface that is connected to the primary ISP.

Make sure to exclude traffic destined to internal servers/IP addresses from PBF. Define a negate rule so that traffic destined to internal IP addresses is not routed through the egress interface defined in the PBF rule.

1. Select **Policies > Policy Based Forwarding** and click **Add.**
2. Give the rule a descriptive **Name** in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination/Application/Service** tab, set the following:
   a. In the Destination Address section, **Add** the IP addresses or address range for servers on the internal network or create an address object for your internal servers. Select **Negate** to exclude the IP addresses or address object listed above from using this rule.
   b. In the Service section, **Add** the **service-http** and **service-https** services to allow HTTP and HTTPS traffic to use the default ports. For all other traffic that is allowed by security policy, the default route will be used.

   To forward all traffic using PBF, set the Service to **Any**.

5. In the **Forwarding** tab, specify the interface to which you want to forward traffic and enable path monitoring.
   a. To forward traffic, set the **Action** to **Forward**, and select the **Egress Interface** and specify the **Next Hop**. In this example, the egress interface is ethernet1/1, and the next hop IP address is 1.1.1.1.

© Palo Alto Networks, Inc.

**PBF for Outbound Access with Dual ISPs**

b. Enable **Monitor** and attach the default monitoring profile, to trigger a failover to the backup ISP. In this example, we do not specify a target IP address to monitor. The firewall will monitor the next hop IP address; if this IP address is unreachable the firewall will direct traffic to the default route specified on the virtual router.

c. (Required if you have asymmetric routes). Select **Enforce Symmetric Return** to ensure that return traffic from the Trust zone to the Internet is forwarded out on the same interface through which traffic ingressed from the Internet.

NAT ensures that the traffic from the Internet is returned to the correct interface/IP address on the firewall.

d. Click **OK** to save the changes.

**PBF for Outbound Access with Dual ISPs**

Step 4 Create NAT rules based on the egress interface and ISP. These rules ensure that the correct source IP address is used for outbound connections.

1. Select **Policies > NAT** and click **Add**.
2. In this example, the NAT rule we create for each ISP is as follows:

**NAT for Primary ISP**

In the **Original Packet** tab,

- **Source Zone**: Trust
- **Destination Zone**: ISP-West

In the **Translated Packet** tab, under Source Address Translation

- **Translation Type**: Dynamic IP and Port
- **Address Type**: Interface Address
- **Interface**: ethernet1/1
- **IP Address**: 1.1.1.2/30

**NAT for Backup ISP**

In the **Original Packet** tab,

- **Source Zone**: Trust
- **Destination Zone**: ISP-East

In the **Translated Packet** tab, under Source Address Translation

- **Translation Type**: Dynamic IP and Port
- **Address Type**: Interface Address
- **Interface**: ethernet1/3
- **IP Address**: 2.2.2.2/30

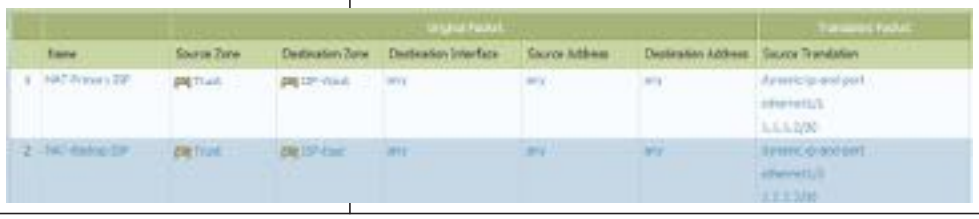

 © Palo Alto Networks, Inc.

| PBF for Outbound Access with Dual ISPs | | |
|---|---|---|
| Step 5 | Create security policy to allow outbound access to the Internet. | To safely enable applications, create a simple rule that allows access to the Internet and attach the security profiles available on the firewall.<br>1. Select **Policies > Security** and click **Add**.<br>2. Give the rule a descriptive **Name** in the **General** tab.<br>3. In the **Source** tab, set the **Source Zone** to Trust.<br>4. In the **Destination** tab, Set the **Destination Zone** to ISP-East and ISP-West.<br>5. In the **Service/ URL Category** tab, leave the default **application-default**.<br>6. In the **Actions** tab, complete these tasks:<br>a. Set the **Action Setting** to **Allow**.<br>b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection and URL filtering, under **Profile Setting.**<br>7. Under **Options**, verify that logging is enabled at the end of a session. Only traffic that matches a security rule is logged.<br> 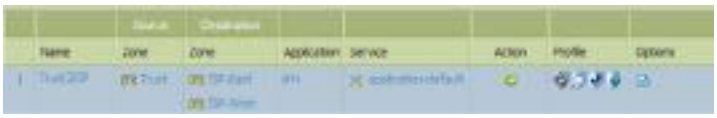 |
| Step 6 | Save the policies to the running configuration on the device. | Click **Commit**. |

**PBF for Outbound Access with Dual ISPs**

**Step 7** Verify that the PBF rule is active and that the primary ISP is used for Internet access.

1. Launch a web browser and access a web server. On the firewall check the traffic log for web-browsing activity.

2. From a client on the network, use the ping utility to verify connectivity to a web server on the Internet. and check the traffic log on the firewall.

```
C:\Users\pm-user1>ping 4.2.2.1
Pinging 4.2.2.1 with 32 bytes of data:
Reply from 4.2.2.1: bytes=32 time=34ms TTL=117
Reply from 4.2.2.1: bytes=32 time=13ms TTL=117
Reply from 4.2.2.1: bytes=32 time=25ms TTL=117
Reply from 4.2.2.1: bytes=32 time=3ms TTL=117
Ping statistics for 4.2.2.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 3ms, Maximum = 34ms, Average = 18ms
```

3. To confirm that the PBF rule is active, use the CLI command **show pbf rule all**

```
admin@PA-NGFW> show pbf rule all
Rule       ID    Rule State Action   Egress IF/VSYS  NextHop
========== === ========== ====== ============== =======
Use ISP-Pr 1     Active     Forward ethernet1/1     1.1.1.1
```

**Step 8** Verify that the failover to the backup ISP occurs and that the Source NAT is correctly applied.

1. Unplug the connection to the primary ISP.
2. Confirm that the PBF rule is inactive with the CLI command **show pbf rule all**

```
admin@PA-NGFW> show pbf rule all
Rule       ID    Rule State Action   Egress IF/VSYS  NextHop
========== === ========== ====== ============== =======
Use ISP-Pr 1   Disabled   Forward   ethernet1/1     1.1.1.1
```

3. Access a web server, and check the traffic log to verify that traffic is being forwarded through the backup ISP.

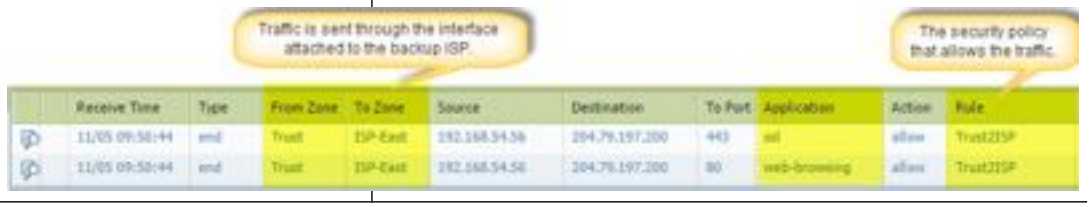


 © Palo Alto Networks, Inc.

**PBF for Outbound Access with Dual ISPs**

4. View the session details to confirm that the NAT rule is working properly.

```
admin@PA-NGFW> show session all
--------------------------------------------------------------
ID Application    State   Type Flag Src[Sport]/Zone/Proto
  (translated IP[Port]) Vsys Dst[Dport]/Zone (translated
  IP[Port])
--------------------------------------------------------------
87212 ssl ACTIVE  FLOW  NS   192.168.54.56[53236]/Trust/6
  (2.2.2.2[12896]) vsys1 204.79.197.200[443]/ISP-East
  (204.79.197.200[443])
```

5. Obtain the session identification number from the output and view the session details. Note that the PBF rule is not used and hence is not listed in the output.

```
admin@PA-NGFW> show session id 87212

Session           87212

        c2s flow:
                source:      192.168.54.56 [Trust]
                dst:         204.79.197.200
                proto:       6
                sport:       53236          dport:      443
                state:       ACTIVE         type:       FLOW
                src user:    unknown
                dst user:    unknown

        s2c flow:
                source:      204.79.197.200 [ISP-East]
                dst:         2.2.2.2
                proto:       6
                sport:       443            dport:      12896
                state:       ACTIVE         type:       FLOW
                src user:    unknown
                dst user:    unknown
start time                      : Wed Nov5 11:16:10 2014
        timeout                         : 1800 sec
        time to live                    : 1757 sec
        total byte count(c2s)           : 1918
        total byte count(s2c)           : 4333
        layer7 packet count(c2s)        : 10
        layer7 packet count(s2c)        : 7
        vsys                            : vsys1
        application                     : ssl
        rule                            : Trust2ISP
        session to be logged at end     : True
        session in session ager         : True
        session synced from HA peer     : False
        address/port translation        : source
       nat-rule                       : NAT-Backup ISP(vsys1)
        layer7 processing               : enabled
        URL filtering enabled           : True
        URL category                    : search-engines
        session via syn-cookies         : False
        session terminated on host      : False
        session traverses tunnel        : False
        captive portal session          : False
        ingress interface               : ethernet1/2
        egress interface                : ethernet1/3
        session QoS rule                : N/A (class 4)
```

# DoS Protection Against Flooding of New Sessions

The following topics describe how to configure DoS protection to better block IP addresses in order to handle high-volume attacks more efficiently.


- DoS Protection Against Flooding of New Sessions
- Configure DoS Protection Against Flooding of New Sessions
- Use the CLI to End a Single Attacking Session
- Identify Sessions That Use an Excessive Percentage of the Packet Buffer
- Discard a Session Without a Commit


 © Palo Alto Networks, Inc.

# DoS Protection Against Flooding of New Sessions

DoS protection against flooding of new sessions is beneficial against high-volume single-session and multiple-session attacks. In a single-session attack, an attacker uses a single session to target a device behind the firewall. If a Security policy rule allows the traffic, the session is established and the attacker initiates an attack by sending packets at a very high rate with the same source IP address and port number, destination IP address and port number, and protocol, trying to overwhelm the target. In a multiple-session attack, an attacker uses multiple sessions (or connections per second [cps]) from a single host to launch a DoS attack.

This feature defends only against DoS attacks of new sessions, that is, traffic that has not been offloaded to hardware. An offloaded attack is not protected by this feature. However, this topic describes how you can create a Security policy rule to reset the client; the attacker reinitiates the attack with numerous connections per second and is blocked by the defenses illustrated in this topic.

▲ Multiple-Session DoS Attack
▲ Single-Session DoS Attack

## Multiple-Session DoS Attack

Configure DoS Protection Against Flooding of New Sessions by configuring a DoS Protection policy rule, which determines the criteria that, when matched by incoming packets, trigger the protect action. The DoS Protection profile counts each new connection toward the Alarm Rate, Activate Rate, and Max Rate thresholds. When the incoming new connections per second exceed the Max Rate allowed, the firewall takes the action specified in the DoS Protection policy rule.

The following figure and table describe how the Security policy rules, DoS Protection policy rules and profile work together in an example.

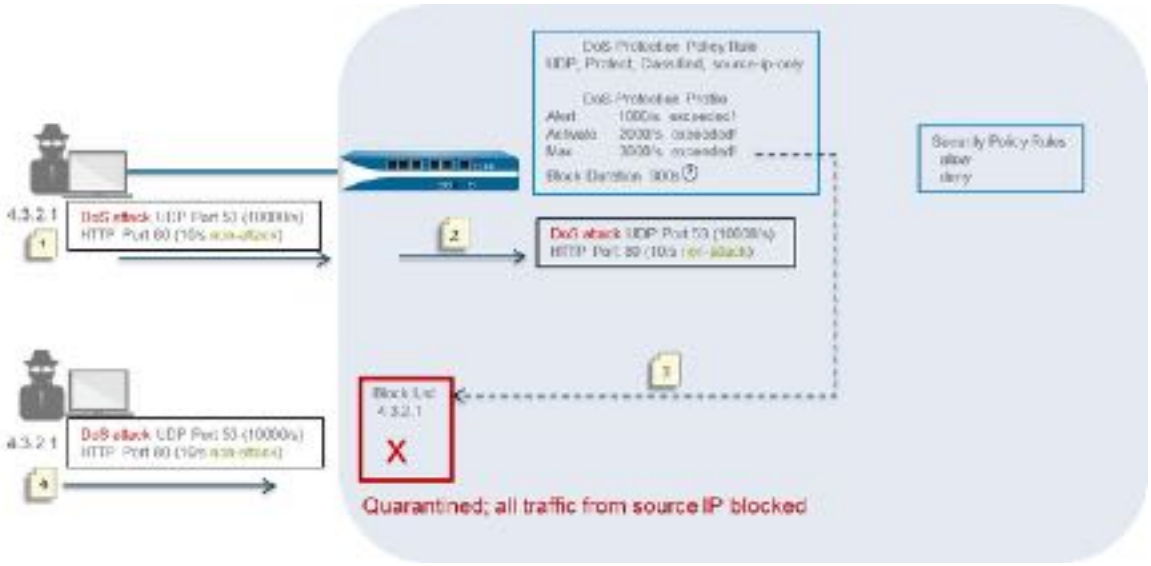

| Sequence of Events as Firewall Quarantines an IP Address | |
|---|---|
| 1 | In this example, an attacker launches a DoS attack at a rate of 10,000 new connections per second to UDP port 53. The attacker also sends 10 new connections per second to HTTP port 80. |
| 2 | The new connections match criteria in the DoS Protection policy rule, such as a source zone or interface, source IP address, destination zone or interface, destination IP address, or a service, among other settings. In this example, the policy rule specifies UDP.<br>The DoS rule also specifies the Protect action and Classified, two settings that dynamically put the DoS Protection Profile settings into effect. The DoS Protection Profile specifies that a Max Rate of 3000 packets per second is allowed. When incoming packets match the DoS rule, new connections per second are counted toward the Alert, Activate, and Max Rate thresholds.<br>You can also use a Security policy rule to block all traffic from the source IP address if you deem that address to be malicious all the time. |
| 3 | The 10,000 new connections per second exceed the Max Rate threshold. When all of the following occur:<br>• the threshold is exceeded,<br>• a Block Duration is specified, and<br>• Classified is set to includes source IP address,<br>the firewall puts the offending source IP address on the block list. |
| 4 | An IP address on the block list is in quarantine, meaning all traffic from that IP address is blocked. The firewall blocks the offending source IP address before additional attack packets reach the security policy. |

The following figure describes in more detail what happens after an IP address that matches the DoS Protection policy rule is put on the block list. It also describes the Block Duration timer.

 © Palo Alto Networks, Inc.

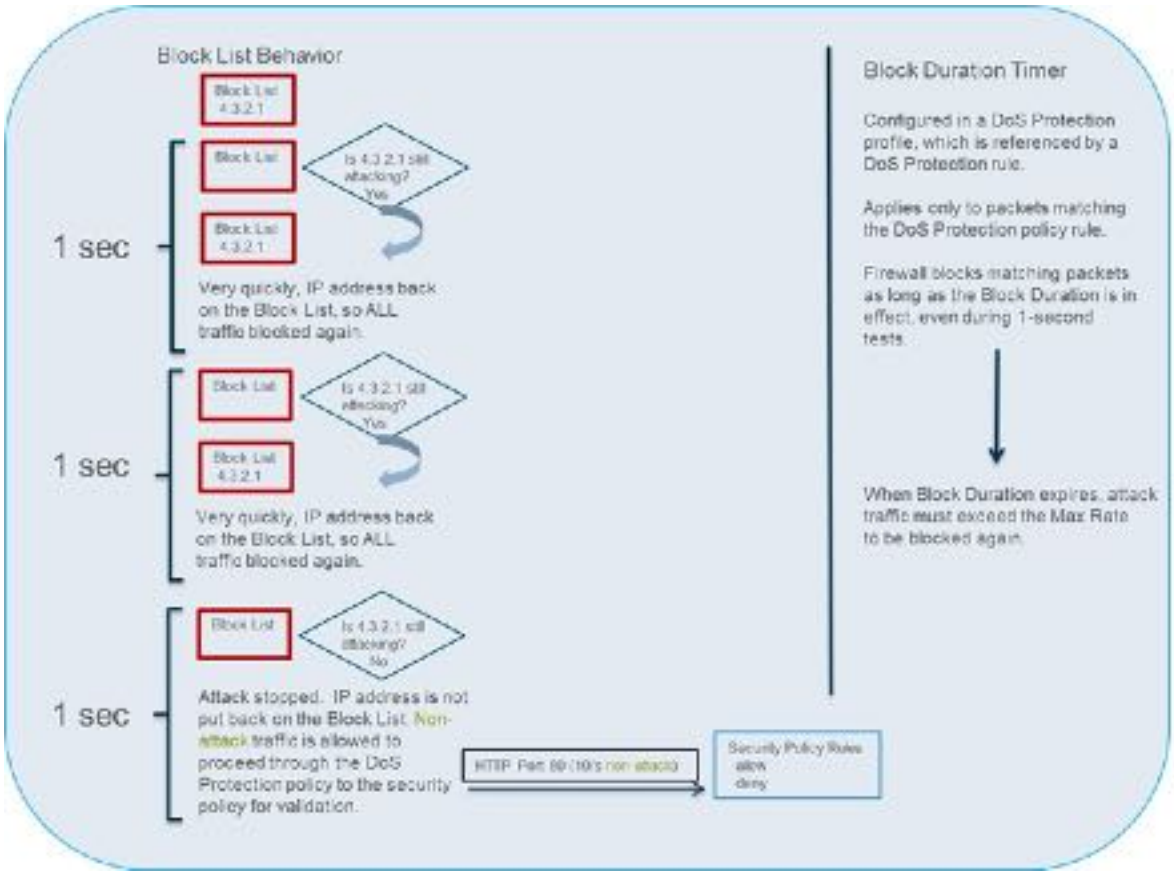


Every one second, the firewall allows the IP address to come off the Block List so that the firewall can test the traffic patterns and determine if the attack is ongoing. The firewall takes the following action:

- During this one-second test period, the firewall allows packets that do not match the DoS Protection policy criteria (HTTP traffic in this example) through the DoS Protection policy rules to the Security policy for validation. Very few packets, if any, have time to get through because the first attack packet that the firewall receives after the IP address is let off the Block List will match the DoS Protection policy criteria, quickly causing the IP address to be placed back on the block list for another second. The firewall repeats this test each second until the attack stops.
- The firewall blocks all attack traffic from going past the DoS Protection policy rules until the Block Duration expires.

When the attack stops, the firewall does not put the IP address back on the Block List. The firewall allows non-attack traffic to proceed through the DoS Protection policy rules to the Security policy rules for validation. You must configure a Security policy rule because without one, an implicit deny rule denies all traffic.

The block list is based on a source zone and source address combination. This behavior allows duplicate IP addresses to exist as long as they are in different zones belonging to separate virtual routers.

The Block Duration setting in a DoS Protection profile specifies how long the firewall blocks the [offending] packets that exactly match a DoS Protection policy rule. The attack traffic remains blocked until the Block Duration expires, after which the attack traffic must again exceed the Max Rate threshold to be blocked again.

If the attacker uses multiple sessions or bots that initiate multiple attack sessions, the sessions count toward the thresholds in the DoS Protection profile without a Security policy deny rule in place. Hence, a single-session attack requires a Security policy deny rule in order for each packet to count toward the thresholds; a multiple-session attack does not.

Therefore, the DoS protection against flooding of new sessions allows the firewall to efficiently defend against a source IP address while attack traffic is ongoing and to permit non-attack traffic to pass as soon as the attack stops. Putting the offending IP address on the block list allows the DoS protection functionality to take advantage of the block list, which is designed to quarantine all activity. Quarantining the IP address from all activity protects against a modern attacker who attempts a rotating application attack, in which the attacker simply changes applications to start a new attack or uses a combination of different attacks in a hybrid DoS attack.

Beginning with PAN-OS 6.1.7, it is a change in behavior that the firewall places the attacking source IP address on the block list. When the attack stops, non-attack traffic is allowed to proceed to the Security policy rules. The attack traffic that matched the DoS Protection profile and DoS Protection policy rules remains blocked until the Block Duration expires.

## Single-Session DoS Attack

A single-session DoS attack typically will not trigger Zone or DoS Protection profiles because they are attacks that are formed after the session is created. These attacks are allowed by the Security policy because a session is allowed to be created, and after the session is created, the attack drives up the packet volume and takes down the target device.

Configure DoS Protection Against Flooding of New Sessions to protect against flooding of new sessions (single-session and multiple-session flooding). In the event of a single-session attack that is underway, additionally Use the CLI to End a Single Attacking Session.

 © Palo Alto Networks, Inc.

# Configure DoS Protection Against Flooding of New Sessions

| Configure DoS Protection Against Flooding of New Sessions | | |
|---|---|---|
| Step 1 | (Required for single-session attack mitigation or attacks that have not triggered the DoS Protection policy threshold; optional for multiple-session attack mitigation)<br><br>Configure Security policy rules to deny traffic from the attacker's IP address and allow other traffic based on your network needs. You can specify any of the match criteria in a Security policy rule, such as source IP address.<br><br> This step is one of the steps typically performed to stop an existing attack. See Use the CLI to End a Single Attacking Session. | • Set Up Basic Security Policies<br>• Components of a Security Policy Rule |

**Configure DoS Protection Against Flooding of New Sessions (Continued)**

Step 2 Configure a DoS Protection profile for flood protection.

Because flood attacks can occur over multiple protocols, as a best practice, activate protection for all of the flood types in the DoS Protection profile.

1. Select **Objects > Security Profiles > DoS Protection** and **Add** a profile **Name**.
2. Select **Classified** as the **Type**.
3. For **Flood Protection**, select the check boxes for all of the following types of flood protection:
   - **SYN Flood**
   - **UDP Flood**
   - **ICMP Flood**
   - **ICMPv6 Flood**
   - **Other IP Flood**
4. (Optional) On each of the flood tabs, change the following thresholds to suit your environment:
   - **Alarm Rate (packets/s)**—Specify the threshold rate (packets per second [pps]) above which a DoS alarm is generated. (Range is 0-2000000; default is 10000.)
   - **Activate Rate (packets/s)**—Specify the threshold rate (pps) above which a DoS response is activated. The DoS response is configured in the **Action** field of the DoS policy where this profile is referenced. When the **Activate Rate** threshold is reached, **Random Early Drop** occurs. (Range is 0-2000000; default is 10000.)
   - **Max Rate (packets/s)**—Specify the threshold rate of incoming packets per second that the firewall allows. When the threshold is exceeded, new packets that arrive are dropped and the **Action** specified in the DoS Policy rule is triggered. (Range is 2-2000000; default is 40000.)

   The default threshold values in this step are only starting points and might not be appropriate for your network. You must analyze the behavior of your network to properly set initial threshold values.
5. On each of the flood tabs, specify the **Block Duration** (in seconds), which is the length of time the firewall blocks packets that match the DoS Protection policy rule that references this profile. Specify a value greater than zero. (Range is 1-21600; default is 300.)

   Set a low Block Duration value if you are concerned that packets you incorrectly identified as attack traffic will be blocked unnecessarily.

   Set a high Block Duration value if you are more concerned about blocking volumetric attacks than you are about incorrectly blocking packets that are not part of an attack.
6. Click **OK**.


© Palo Alto Networks, Inc.

| Configure DoS Protection Against Flooding of New Sessions (Continued) | | |
|---|---|---|
| Step 3 | Configure a DoS Protection policy rule that specifies the criteria for matching the incoming traffic. | 1. Select **Policies > DoS Protection** and **Add** a **Name** on the **General** tab. The name is case-sensitive and can be a maximum of 31 characters, including letters, numbers, spaces, hyphens, and underscores.<br>2. On the **Source** tab, choose the **Type** to be a **Zone** or **Interface**, and then **Add** the zone(s) or interface(s).<br>3. (Optional) For **Source Address**, select **Any** for any incoming IP address to match the rule or **Add** an address object such as a geographical region.<br>4. (Optional) For **Source User**, select **any** or specify a user.<br>5. (Optional) Select **Negate** to match any sources except those you specify.<br>6. (Optional) On the **Destination** tab, choose the **Type** to be a **Zone** or **Interface**, and then **Add** the destination zone(s) or interface(s). For example, enter the security zone you want to protect.<br>7. (Optional) For **Destination Address**, select **Any** or enter the IP address of the device you want to protect.<br>8. (Optional) On the **Option/Protection** tab, **Add** a **Service**. Select a service or click **Service** and enter a **Name**. Select **TCP** or **UDP**. Enter a **Destination Port**. Not specifying a particular service allows the rule to match a flood of any protocol type without regard to an application-specific port.<br>9. On the **Option/Protection** tab, for **Action**, select **Protect**.<br>10. Select the **Classified** check box.<br>11. For **Profile**, select the name of the **DoS Protection** profile you created in Step 2.<br>12. For **Address**, select **source-ip-only** or **src-dest-ip-both**, which determines the type of IP address to which the rule applies. Choose the setting based on how you want the firewall to identify offending traffic.<br>• Specify **source-ip-only** if you want the firewall to classify only on the source IP address. Because attackers often test the entire network for hosts to attack, **source-ip-only** is the typical setting for a wider examination.<br>• Specify **src-dest-ip-both** if you want to protect only against DoS attacks on the server that has a specific destination address and also ensure that every source IP address will not surpass a specific connections-per-second threshold to that server.<br>13. Click **OK**. |
| Step 4 | Save the configuration. | Click **Commit**. |

## Use the CLI to End a Single Attacking Session

To mitigate a single-session DoS attack, you would still Configure DoS Protection Against Flooding of New Sessions in advance. At some point after you configure the feature, a session might be established before you realize a DoS attack (from the IP address of that session) is underway. When you see a single-session DoS attack, perform the following task to end the session, so that subsequent connection attempts from that IP address trigger the DoS protection against flooding of new sessions.

| Use the CLI to End a Single Attacking Session | |
|---|---|
| Step 1 | Identify the source IP address that is causing the attack.<br>For example, use the firewall Packet Capture feature with a destination filter to collect a sample of the traffic going to the destination IP address. Alternatively, in PAN-OS 7.0 and later, you can use ACC to filter on destination address to view the activity to the target host being attacked. |
| Step 2 | Create a DoS Protection policy rule that will block the attacker's IP address after the attack thresholds are exceeded. |
| Step 3 | Create a Security policy rule to deny the source IP address and its attack traffic. |
| Step 4 | End any existing attacks from the attacking source IP address by executing the `clear session all filter source <ip-address>` operational command.<br>Alternatively, if you know the session ID, you can execute the `clear session id <value>` command to end that session only.<br>If you use the `clear session all filter source <ip-address>` command, all sessions matching the source IP address are discarded, which can include both good and bad sessions. |

After you end the existing attack session, any subsequent attempts to form an attack session are blocked by the Security policy. The DoS Protection policy counts all connection attempts toward the thresholds. When the Max Rate threshold is exceeded, the source IP address is blocked for the Block Duration, as described in Sequence of Events as Firewall Quarantines an IP Address.

 © Palo Alto Networks, Inc.

## Identify Sessions That Use an Excessive Percentage of the Packet Buffer

When a firewall exhibits signs of resource depletion, it might be experiencing an attack that is sending an overwhelming number of packets. In such events, the firewall starts buffering inbound packets. With PAN-OS 6.1.10 and later PAN-OS 6.1 releases, you can quickly identify the sessions that are using an excessive percentage of the packet buffer and mitigate their impact by discarding them.

Perform the following task on any hardware-based firewall platform (not a VM-Series firewall) to identify, for each slot and dataplane, the packet buffer percentage used, the top five sessions using more than two percent of the packet buffer, and the source IP addresses associated with those sessions. Having that information allows you to take appropriate action.

**View Firewall Resource Usage, Top Sessions, and Session Details**

Step 1 View firewall resource usage, top sessions, and session details. Execute the following operational command in the CLI (sample output from the command follows):

```
admin@PA-7050> show running resource-monitor ingress-backlogs

-- SLOT:s1, DP:dp1 --
USAGE - ATOMIC: 92%  TOTAL: 93%

TOP SESSIONS:
SESS-ID      PCT   GRP-ID   COUNT
6            92%   1        156
                   7        1732

SESSION DETAILS
SESS-ID PROTO SZONE SRC           SPORT  DST       DPORT  IGR-IF    EGR-IF       APP
6       6     trust 192.168.2.35 55653  10.1.8.89 80  ethernet1/21 ethernet1/22
undecided
```

The command displays a maximum of the top five sessions that each use 2% or more of the packet buffer.

The sample output above indicates that Session 6 is using 92% of the packet buffers with TCP packets (protocol 6) coming from source IP address 192.168.2.35.

- **GRP-ID**—indicates an internal stage of processing packets.
- **COUNT**—indicates how many packets are in that GRP-ID for that session.
- **SESS-ID**—indicates the global session ID that is used in all other `show session` commands. The global session ID is unique within the firewall.
- **APP**—indicates the application name extracted from the Session information, which can help you determine whether the traffic is legitimate. For example, spoofed packets that use a common TCP or UDP port but indicate an APP of `undecided` are possibly attack traffic.

To restrict the display output:

On a PA-7000 Series platform, you can limit output to a slot, a dataplane, or both. For example:

```
admin@PA-7050> show running resource-monitor ingress-backlogs slot s1
admin@PA-7050> show running resource-monitor ingress-backlogs slot s1 dp dp1
```

On a PA-5000 Series platform, you can limit output to a dataplane. For example:

```
admin@PA-5060> show running resource-monitor ingress-backlogs dp dp1
```

**View Firewall Resource Usage, Top Sessions, and Session Details**

Step 2 Use the command output to determine whether the source at the source IP address using a high percentage of the packet buffer is sending legitimate or attack traffic.

In the sample output above, a single-session attack is likely occurring. A single session (Session ID 6) is using 92% of the packet buffer for Slot 1, DP 1, and the application at that point is `undecided`.

- If you determine a single user is sending an attack and the traffic is in the slow path, you can Use the CLI to End a Single Attacking Session. At a minimum, you can Configure DoS Protection Against Flooding of New Sessions.
- If the traffic is offloaded to hardware (the traffic is in the fast path), clearing the session does not help because the software then has to handle the barrage of packets. You should Discard a Session Without a Commit instead.

To see whether a session is offloaded or not, use the `show session id <session-id>` operational command in the CLI as shown in the following example. The `layer7 processing value indicates completed` for sessions offloaded or `enabled` for sessions not offloaded.

```
admin@PA-5060> show session id 68088184

Session       68088184

        c2s flow:
                source:      1.1.42.15 [trust]
                dst:         1.2.27.99
                proto:       6
                sport:       55993           dport:      6881
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
                offload:     Yes

        s2c flow:
                source:      1.2.27.99 [untrust]
                dst:         1.1.42.15
                proto:       6
                sport:       6881            dport:      55993
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
                offload:     Yes

        DP                                   : 2
        index(local):                        : 979328
        start time                           : Tue Oct 27 14:20:09 2015
        timeout                              : 1200 sec
        time to live                         : 1167 sec
        total byte count(c2s)                : 270
        total byte count(s2c)                : 270
        layer7 packet count(c2s)             : 3
        layer7 packet count(s2c)             : 3
        vsys                                 : vsys1
        application                          : bittorrent
        rule                                 : rule1
        session to be logged at end          : True
        session in session ager              : True
        session updated by HA peer           : False
        layer7 processing                    : completed
        URL filtering enabled                : False
        session via syn-cookies              : False
        session terminated on host           : False
        session traverses tunnel             : False
        captive portal session               : False
        ingress interface                    : ethernet1/21
        egress interface                     : ethernet1/22
        session QoS rule                     : N/A (class 4)
        tracker stage l7proc                 : ctd decoder bypass
        end-reason                           : unknown
```

 © Palo Alto Networks, Inc.

## Discard a Session Without a Commit

With PAN-OS 6.1.10 and later PAN-OS 6.1 releases, you can perform this task to permanently discard a session, such as a session that is overloading the packet buffer. No commit is required; the session is discarded immediately after executing the command. The commands you use for this task apply to both offloaded and non-offloaded sessions.

| Discard a Session Without a Commit | |
|---|---|
| Step 1 | In the CLI, execute the following operational command on any hardware platform:<br>`admin@PA-7050> request session-discard [timeout <seconds>] [reason <reason-string>] id <session-id>`<br>The default timeout is 3600 seconds. |
| Step 2 | Verify that sessions have been discarded:<br>`admin@PA-7050> show session all filter state discard` |

 © Palo Alto Networks, Inc.

# Virtual Systems

This topic describes virtual systems, their benefits, typical use cases, and how to configure them. It also provides links to other topics where virtual systems are documented as they function with other features.


- Virtual Systems Overview
- Communication Between Virtual Systems
- Shared Gateway
- Configure Virtual Systems
- Configure Inter-Virtual System Communication within the Firewall
- Configure a Shared Gateway
- Virtual System Functionality with Other Features


© Palo Alto Networks, Inc.

# Virtual Systems Overview

Virtual systems are separate, logical firewall instances within a single physical Palo Alto Networks firewall. Rather than using multiple firewalls, managed service providers and enterprises can use a single pair of firewalls (for high availability) and enable virtual systems on them. Each virtual system (vsys) is an independent, separately-managed firewall with its traffic kept separate from the traffic of other virtual systems.

This topic includes the following:


- Virtual System Components and Segmentation
- Benefits of Virtual Systems
- Use Cases for Virtual Systems
- Platform Support and Licensing for Virtual Systems
- Restrictions on Virtual Systems
- Administrative Roles for Virtual Systems
- Shared Objects for Virtual Systems


## Virtual System Components and Segmentation

A virtual system is an object that creates an administrative boundary, as shown in the following figure.

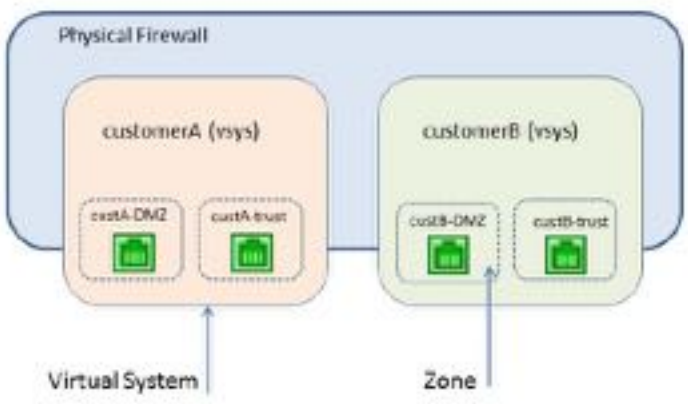


A virtual system consists of a set of physical and logical interfaces and subinterfaces (including VLANs and virtual wires), virtual routers, and security zones. You choose the deployment mode(s) (any combination of virtual wire, Layer 2, or Layer 3) of each virtual system. By using virtual systems, you can segment any of the following:

- Administrative access
- The management of all policies (security, NAT, QoS, policy-based forwarding, decryption, application override, captive portal, and DoS protection)
- All objects (such as address objects, application groups and filters, dynamic block lists, security profiles, decryption profiles, custom objects, etc.)

 © Palo Alto Networks, Inc.

- User-ID
- Certificate management
- Server profiles
- Logging, reporting, and visibility functions

Virtual systems affect the security functions of the firewall, but virtual systems alone do not affect networking functions such as static and dynamic routing. You can segment routing for each virtual system by creating one or more virtual routers for each virtual system, as in the following use cases:

- If you have virtual systems for departments of one organization, and the network traffic for all of the departments is within a common network, you can create a single virtual router for multiple virtual systems.
- If you want routing segmentation and each virtual system's traffic must be isolated from other virtual systems, you can create one or more virtual routers for each virtual system.

## Benefits of Virtual Systems

Virtual systems provide the same basic functions as a physical firewall, along with additional benefits:

- **Segmented administration**—Different organizations (or customers or business units) can control (and monitor) a separate firewall instance, so that they have control over their own traffic without interfering with the traffic or policies of another firewall instance on the same physical device.
- **Scalability**—After the physical firewall is configured, adding or removing customers or business units can be done efficiently. An ISP, managed security service provider, or enterprise can provide different security services to each customer.
- **Reduced capital and operational expenses**—Virtual systems eliminate the need to have multiple physical firewalls at one location because virtual systems co-exist on one firewall. By not having to purchase multiple firewalls, an organization can save on the hardware expense, electric bills, and rack space, and can reduce maintenance and management expenses.

## Use Cases for Virtual Systems

There are many ways to use virtual systems in a network. One common use case is for an ISP or a managed security service provider (MSSP) to deliver services to multiple customers with a single firewall. Customers can choose from a wide array of services that can be enabled or disabled easily. The firewall's role-based administration allows the ISP or MSSP to control each customer's access to functionality (such as logging and reporting) while hiding or offering read-only capabilities for other functions.

Another common use case is within a large enterprise that requires different firewall instances because of different technical or confidentiality requirements among multiple departments. Like the above case, different groups can have different levels of access while IT manages the firewall itself. Services can be tracked and/or billed back to departments to thereby make separate financial accountability possible within an organization.

© Palo Alto Networks, Inc.

## Platform Support and Licensing for Virtual Systems

Virtual systems are supported on the PA-2000, PA-3000, PA-4000, PA-5000, and PA-7050 Series firewalls. Each firewall series supports a base number of virtual systems; the number varies by platform. A Virtual Systems license is required in the following cases:

- To support multiple virtual systems on PA-2000 or PA-3000 Series firewalls.
- To create more than the base number of virtual systems supported on a platform.

For license information, see Activate Licenses and Subscriptions. For the base and maximum number of virtual systems supported, see Compare Firewalls tool.

Multiple virtual systems are not supported on the PA-200, PA-500 or VM-Series firewalls.

## Restrictions on Virtual Systems

- Within a Palo Alto Networks firewall, a packet may hop from one virtual system to another virtual system or a shared gateway. A packet may not traverse more than two virtual systems or shared gateways. For example, a packet cannot go from one virtual system to a shared gateway to a second virtual system within the firewall.
- Only one Shared Gateway is permitted on a firewall.

## Administrative Roles for Virtual Systems

A **superuser** administrator can create virtual systems and add a **Device Administrator**, **vsysadmin**, or **vsysreader**. A **Device Administrator** can access all virtual systems, but cannot add administrators. The two types of virtual system administrative roles are:

- **vsysadmin**—Grants full access to a virtual system.
- **vsysreader**—Grants read-only access to a virtual system.

A virtual system administrator can view logs of only the virtual systems assigned to that administrator. Someone with **superuser** or **Device Admin** permission can view all of the logs or select a virtual system to view.

Persons with **vsysadmin** permission can commit configurations for only the virtual systems assigned to them.

## Shared Objects for Virtual Systems

If your administrator account extends to multiple virtual systems, you can choose to configure objects (such as an address object) and policies for a specific virtual system or as shared objects, which apply to all of the virtual systems on the firewall. If you try to create a shared object with the same name and type as an existing object in a virtual system, the virtual system object is used.

 © Palo Alto Networks, Inc.

# Communication Between Virtual Systems

There are two typical scenarios where communication between virtual systems (inter-vsys traffic) is desirable. In a multi-tenancy environment, communication between virtual systems can occur by having traffic leave the firewall, go through the Internet, and re-enter the firewall. In a single organization environment, communication between virtual systems can remain within the firewall. This section discusses both scenarios.


- Inter-VSYS Traffic That Must Leave the Firewall
- Inter-VSYS Traffic That Remains Within the Firewall
- Inter-VSYS Communication Uses Two Sessions

## Inter-VSYS Traffic That Must Leave the Firewall

An ISP that has multiple customers on a firewall (known as multi-tenancy) can use a virtual system for each customer, and thereby give each customer control over its virtual system configuration. The ISP grants **vsysadmin** permission to customers. Each customer's traffic and management are isolated from the others. Each virtual system must be configured with its own IP address and one or more virtual routers in order to manage traffic and its own connection to the Internet.

If the virtual systems need to communicate with each other, that traffic goes out the firewall to another Layer 3 routing device and back to the firewall, even though the virtual systems exist on the same physical firewall, as shown in the following figure.

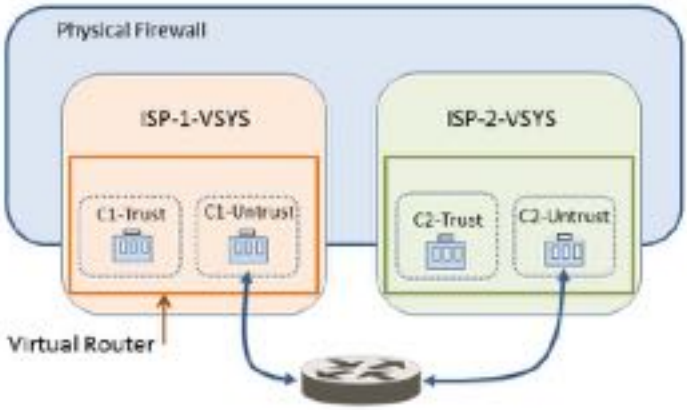


 © Palo Alto Networks, Inc.

# Inter-VSYS Traffic That Remains Within the Firewall

Unlike the preceding multi-tenancy scenario, virtual systems on a firewall can be under the control of a single organization. The organization wants to both isolate traffic between virtual systems and allow communications between virtual systems. This common use case arises when the organization wants to provide departmental separation and still have the departments be able to communicate with each other or connect to the same network(s). In this scenario, the inter-vsys traffic remains within the firewall, as described in the following topics:


- External Zone
- External Zones and Security Policies For Traffic Within a Firewall


## External Zone

The communication desired in the use case above is achieved by configuring security policies that point to or from an *external* zone. An external zone is a security object that is associated with a specific virtual system that it can reach; the zone is external to the virtual system. A virtual system can have only one external zone, regardless of how many security zones the virtual system has within it. External zones are required to allow traffic between zones in different virtual systems, without the traffic leaving the firewall.

The virtual system administrator configures the security policies needed to allow traffic between two virtual systems. Unlike security zones, an external zone is not associated with an interface; it is associated with a virtual system. The security policy allows or denies traffic between the security (internal) zone and the external zone.

Because external zones do not have interfaces or IP addresses associated with them, some zone protection profiles are not supported on external zones.

Remember that each virtual system is a separate instance of a firewall, which means that each packet moving between virtual systems is inspected for security policy and App-ID evaluation.

## External Zones and Security Policies For Traffic Within a Firewall

In the following example, an enterprise has two separate administrative groups: the departmentA and departmentB virtual systems. The following figure shows the external zone associated with each virtual system, and traffic flowing from one trust zone, out an external zone, into an external zone of another virtual system, and into its trust zone.

© Palo Alto Networks, Inc.

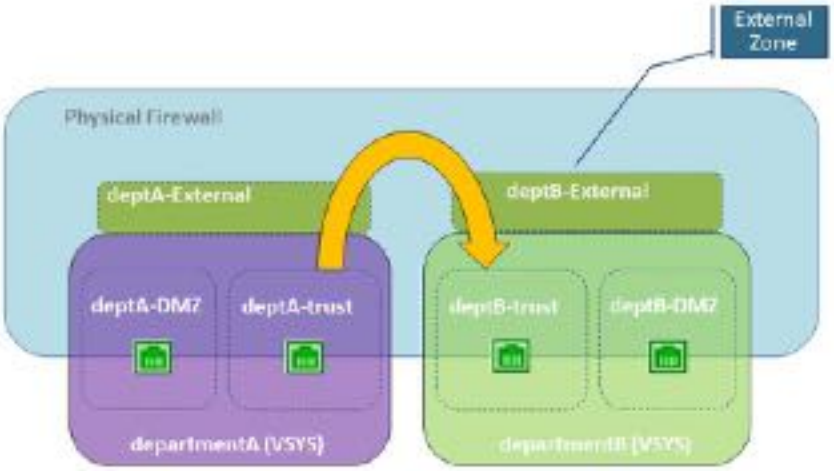


In order to create external zones, the device administrator must configure the virtual systems so that they are *visible* to each other. External zones do not have security policies between them because their virtual systems are visible to each other.

To communicate between virtual systems, the ingress and egress interfaces on the firewall are either assigned to a single virtual router or else they are connected using inter-virtual router static routes. The simpler of these two approaches is to assign all virtual systems that must communicate with each other to a single virtual router.

There might be a reason that the virtual systems need to have their own virtual router, for example, if the virtual systems use overlapping IP address ranges. Traffic can be routed between the virtual systems, but each virtual router must have static routes that point to the other virtual router(s) as the next hop.

Referring to the scenario in the figure above, we have an enterprise with two administrative groups: departmentA and departmentB. The departmentA group manages the local network and the DMZ resources. The departmentB group manages traffic in and out of the sales segment of the network. All traffic is on a local network, so a single virtual router is used. There are two external zones configured for communication between the two virtual systems. The departmentA virtual system has three zones used in security policies: deptA-DMZ, deptA-trust, and deptA-External. The departmentB virtual system also has three zones: deptB-DMZ, deptB-trust, and deptB-External. Both groups can control the traffic passing through their virtual systems.

In order to allow traffic from deptA-trust to deptB-trust, two security policies are required. In the following figure, the two vertical arrows indicate where the security policies (described below the figure) are controlling traffic.

 © Palo Alto Networks, Inc.

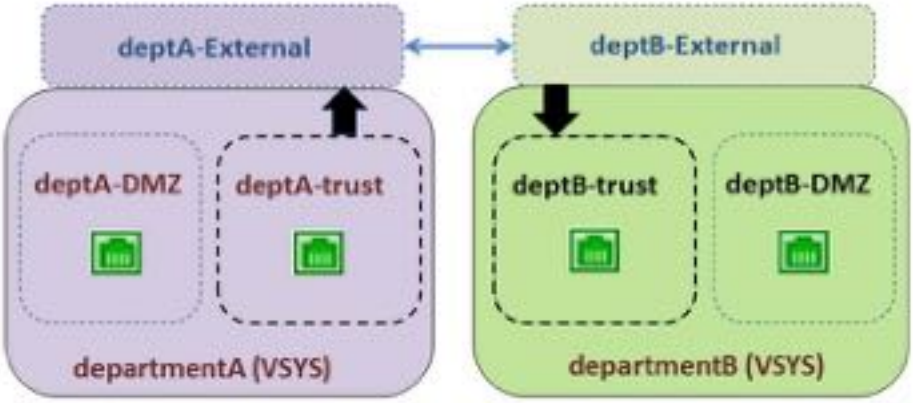


- Security Policy 1: In the preceding figure, traffic is destined for the deptB-trust zone. Traffic leaves the deptA-trust zone and goes to the deptA-External zone. A security policy must allow traffic from the source zone (deptA-trust) to the destination zone (deptA-External). A virtual system allows any policy type to be used for this traffic, including NAT.

  No policy is needed between external zones because traffic sent to an external zone appears in and has automatic access to the other external zones that are visible to the original external zone.

- Security Policy 2: In the preceding figure, the traffic from deptB-External is still destined to the deptB-trust zone, and a security policy must be configured to allow it. The policy must allow traffic from the source zone (deptB-External) to the destination zone (deptB-trust).

The departmentB virtual system could be configured to block traffic from the departmentA virtual system, and vice versa. Like traffic from any other zone, traffic from external zones must be explicitly allowed by policy to reach other zones in a virtual system.

In addition to external zones being required for inter-virtual system traffic that does not leave the firewall, external zones are also required if you configure a Shared Gateway, in which case the traffic is intended to leave the firewall.

## Inter-VSYS Communication Uses Two Sessions

It is helpful to understand that communication between two virtual systems uses two sessions, unlike the one session used for a single virtual system. Let's compare the scenarios.

Scenario 1—Vsys1 has two zones: trust1 and untrust1. A host in the trust1 zone initiates traffic when it needs to communicate with a device in the untrust1 zone. The host sends traffic to the firewall, and the firewall creates a new session for source zone trust1 to destination zone untrust1. Only one session is needed for this traffic.

Scenario 2—A host from vsys1 needs to access a server on vsys2. A host in the trust1 zone initiates traffic to the firewall, and the firewall creates the first session: source zone trust1 to destination zone untrust1. Traffic is routed to vsys2, either internally or externally. Then the firewall creates a second session: source zone untrust2 to destination zone trust2. Two sessions are needed for this inter-vsys traffic.

 © Palo Alto Networks, Inc.

# Shared Gateway

This topic includes the following information about shared gateways:

- External Zones and Shared Gateway
- Networking Considerations for a Shared Gateway

## External Zones and Shared Gateway

A shared gateway is an interface that multiple virtual systems share in order to communicate over the Internet. Each virtual system requires an External Zone, which acts as an intermediary, for configuring security policies that allow or deny traffic from the virtual system's internal zone to the shared gateway.

The shared gateway uses a single virtual router to route traffic for all virtual systems. A shared gateway is used in cases when an interface does not need a full administrative boundary around it, or when multiple virtual systems must share a single Internet connection. This second case arises if an ISP provides an organization with only one IP address (interface), but multiple virtual systems need external communication.

Unlike the behavior between virtual systems, security policy and App-ID evaluations are not performed between a virtual system and a shared gateway. That is why using a shared gateway to access the Internet involves less overhead than creating another virtual system to do so.

In the following figure, three customers share a firewall, but there is only one interface accessible to the Internet. Creating another virtual system would add the overhead of App-ID and security policy evaluation for traffic being sent to the interface through the added virtual system. To avoid adding another virtual system, the solution is to configure a shared gateway, as shown in the following diagram.

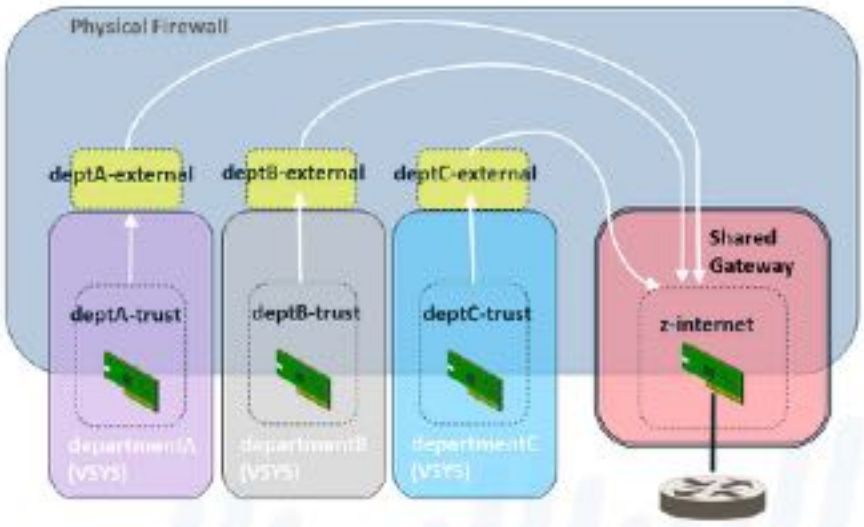


The shared gateway has one globally-routable IP address used to communicate with the outside world. Interfaces in the virtual systems have IP addresses too, but they can be private, non-routable IP addresses.

You will recall that an administrator must specify whether a virtual system is visible to other virtual systems. Unlike a virtual system, a shared gateway is always visible to all of the virtual systems on the firewall.

© Palo Alto Networks, Inc.

A shared gateway ID number appears as **sg<ID>** on the web interface. It is recommended that you name your shared gateway with a name that includes its ID number.

When you add objects such as zones or interfaces to a shared gateway, the shared gateway appears as an available virtual system in the vsys drop-down menu.

A shared gateway is a limited version of a virtual system; it supports NAT, policy-based forwarding (PBF), and DoS policies, but does not support security, QoS, decryption, application override, or captive portal policies.

## Networking Considerations for a Shared Gateway

Keep the following in mind while you are configuring a shared gateway.

- The virtual systems in a shared gateway scenario access the Internet through the shared gateway's physical interface, using a single IP address. If the IP addresses of the virtual systems are not globally routable, configure source NAT to translate those addresses to globally-routable IP addresses.
- A virtual router routes the traffic for all of the virtual systems through the shared gateway.
- The default route for the virtual systems should point to the shared gateway.
- Security policies must be configured for each virtual system to allow the traffic between the internal zone and external zone, which is visible to the shared gateway.
- A device administrator should control the virtual router, so that no member of a virtual system can affect the traffic of other virtual systems.

To save configuration time and effort, consider the following advantages of a shared gateway:

- Rather than configure NAT for multiple virtual systems associated with a shared gateway, you can configure NAT for the shared gateway.
- Rather than configure policy-based routing (PBR) for multiple virtual systems associated with a shared gateway, you can configure PBR for the shared gateway.

 © Palo Alto Networks, Inc.

# Configure Virtual Systems

Creating a virtual system requires that you have the following:

- A **superuser** administrative role.
- An interface configured.
- A Virtual Systems license if configuring a PA-2000 or PA-3000 Series firewall, or if creating more than the base number of virtual systems supported on the platform. See Platform Support and Licensing for Virtual Systems.

Perform the following procedure to create and configure virtual systems.

| Configure a Virtual System | |
|---|---|
| Step 1 Enable virtual systems. | 1. Select **Device > Setup > Management** and edit the **General Settings**.<br>2. Select the **Multi Virtual System Capability** check box and click **OK**. This action triggers a commit if you approve it.<br>Only after completing Step 1 will the **Device** tab display the **Virtual Systems** and **Shared Gateways** options. |
| Step 2 Create a virtual system. | 1. Select **Device > Virtual Systems**, click **Add** and enter a virtual system **ID**, which is appended to "vsys." Range: 1-255.<br>The default ID is 1, which makes the default virtual system **vsys1**. This default appears even on platforms that do not support multiple virtual systems.<br>2. Check the **Allow forwarding of decrypted content** check box if you want to allow the firewall to forward decrypted content to an outside service. For example, you must enable this option for the firewall to be able to send decrypted content to WildFire for analysis.<br>3. Enter a descriptive **Name** for the virtual system. A maximum of 31 alphanumeric, space, and underscore characters is allowed. |

| Configure a Virtual System | | |
|---|---|---|
| Step 3 | Assign interfaces to the virtual system.<br>The virtual routers, vwires, or VLANs can either be configured already or you can configure them later, at which point you specify the virtual system associated with it. The procedure to configure a virtual router, for example, is in Step 6 below. | 1. On the **General** tab, select a **DNS Proxy** profile if you want to apply DNS proxy rules to the interface.<br>2. In the **Interfaces** field, click **Add** to enter the interfaces or subinterfaces to assign to the virtual system. An interface can belong to only one virtual system.<br>3. Do any of the following, based on the deployment type(s) you need in the virtual system:<br>• In the **VLANs** field, click **Add** to enter the VLAN(s) to assign to the vsys.<br>• In the **Virtual Wires** field, click **Add** to enter the virtual wire(s) to assign to the vsys.<br>• In the **Virtual Routers** field, click **Add** to enter the virtual router(s) to assign to the vsys.<br>4. In the **Visible Virtual System** field, check all virtual systems that should be made visible to the virtual system being configured. This is required for virtual systems that need to communicate with each other.<br>In a multi-tenancy scenario where strict administrative boundaries are required, no virtual systems would be checked.<br>5. Click **OK**. |
| Step 4 | (Optional) Limit the resource allocations for sessions, rules, and VPN tunnels allowed for the virtual system. The flexibility of being able to allocate limits per virtual system allows you to effectively control device resources. | 1. On the **Resource** tab, optionally set limits for a virtual system. There are no default values.<br>• **Sessions Limit**—Range: 1-262144<br>• **Security Rules**—Range: 0-2500<br>• **NAT Rules**—Range: 0-3000<br>• **Decryption Rules**—Range: 0-250<br>• **QoS Rules**—Range: 0-1000<br>• **Application Override Rules**—Range: 0-250<br>• **Policy Based Forwarding Rules**—Range: 0-500<br>• **Captive Portal Rules**—Range: 0-1000<br>• **DoS Protection Rules**—Range: 0-1000<br>• **Site to Site VPN Tunnels**—Range: 0-1024<br>• **Concurrent SSL VPN Tunnels**—Range: 0-1024<br>2. Click **OK**. |
| Step 5 | Save the configuration. | Click **Commit** and **OK**. The virtual system is now an object accessible from the **Objects** tab. |

 © Palo Alto Networks, Inc.

| Configure a Virtual System | | |
|---|---|---|
| Step 6 | Create at least one virtual router for the virtual system in order to make the virtual system capable of networking functions, such as static and dynamic routing.<br>Alternatively, your virtual system might use a VLAN or a virtual wire, depending on your deployment. | 1. Select **Network > Virtual Routers** and **Add** a **virtual router** by **Name**.<br>2. **Add** the **Interfaces** that belong to the virtual router.<br>3. Click **OK**. |
| Step 7 | Configure a security zone for each interface in the virtual system. | For at least one interface, create a Layer 3 security zone. See Set Up Interfaces and Zones. |
| Step 8 | Configure the security policies allowing or denying traffic to and from the zones in the virtual system. | See Set Up Basic Security Policies. |
| Step 9 | Save the configuration. | Click **Commit** and **OK**.<br>After a virtual system is created, a device administrator can use the CLI to commit a configuration for only a specific virtual system:<br>`PA-5060> commit partial vsys vsys<id>` |
| Step 10 | (Optional) View the security policies configured for a virtual system. | Open an SSH session to use the CLI. To view the security policies for a virtual system, in operational mode, use the following commands:<br>`PA-5060> set system setting target-vsys <vsys-id>`<br>`PA-5060> show running security-policy` |

# Configure Inter-Virtual System Communication within the Firewall

Perform this task if you have a use case, perhaps within a single enterprise, where you want the virtual systems to be able to communicate with each other within the firewall. Such a scenario is described in Inter-VSYS Traffic That Remains Within the Firewall. This task presumes:

- ❒ You completed the task, Configure Virtual Systems.
- ❒ When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate with each other to be visible to each other.

| **Configure Inter-Virtual System Communication within the Firewall** | | |
|---|---|---|
| Step 1 | Configure an external zone for each virtual system. | 1. Select **Network > Zones** and **Add** a new zone by **Name**.<br>2. For **Location**, select the virtual system for which you are creating an external zone.<br>3. For **Type**, select **External**.<br>4. For **Virtual Systems**, enter the virtual system that the external zone can reach.<br>5. **Zone Protection Profile**—Optionally select a zone protection profile (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.<br>6. **Log Setting—**Optionally select a log forwarding profile for forwarding zone protection logs to an external system.<br>7. Optionally select the **Enable User Identification** check box to enable User-ID for the external zone.<br>8. Click **OK**. |
| Step 2 | Configure the security policies allowing or denying traffic from the internal zones to the external zone of the virtual system, and vice versa. | • See Set Up Basic Security Policies.<br>• See Inter-VSYS Traffic That Remains Within the Firewall. |
| Step 3 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

# Configure a Shared Gateway

Perform this task if you need multiple virtual systems to share an interface (a Shared Gateway) to the Internet. This task presumes:

- You configured an interface with a globally-routable IP address, which will be the shared gateway.
- You completed the prior task, Configure Virtual Systems. For the interface, you chose the external-facing interface with the globally-routable IP address.
- When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate to be visible to each other.

| **Configure a Shared Gateway** | | |
|---|---|---|
| Step 1 | Configure a Shared Gateway. | 1. Select **Device > Shared Gateway** and enter an **ID**.<br>2. Enter a helpful **Name**, preferably including the **ID** of the gateway.<br>3. In the **DNS Proxy** field, select a DNS proxy profile if you want to apply DNS proxy rules to the interface.<br>4. **Add** an **Interface** that connects to the outside world.<br>5. Click **OK**. |
| Step 2 | Configure the zone for the shared gateway.<br> When adding objects such as zones or interfaces to a shared gateway, the shared gateway itself will be listed as an available vsys in the **VSYS** drop-down menu. | 1. Select **Network > Zones** and **Add** a new zone by **Name**.<br>2. For **Location**, select the shared gateway for which you are creating a zone.<br>3. For **Type**, select **Layer3**.<br>4. **Zone Protection Profile**—Optionally select a zone protection profile (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.<br>5. **Log Setting—**Optionally select a log forwarding profile for forwarding zone protection logs to an external system.<br>6. Optionally select the **Enable User Identification** check box to enable User-ID for the shared gateway.<br>7. Click **OK**. |
| Step 3 | Save the configuration. | Click **Commit**. |

# Virtual System Functionality with Other Features

Many of the firewall's features and functionality are capable of being configured, viewed, logged, or reported per virtual system. Therefore, virtual systems are mentioned in other relevant locations in the documentation and that information is not repeated here. Some of the specific chapters are the following:

- If you are configuring Active/Passive HA, the two firewalls must have the same virtual system capability (single or multiple virtual system capability). See High Availability.
- To configure QoS for virtual systems, see Configure QoS for a Virtual System.
- For information about configuring a firewall with virtual systems in a virtual wire deployment that uses subinterfaces (and VLAN tags), see the Networking section, Virtual Wire Subinterfaces.

 © Palo Alto Networks, Inc.

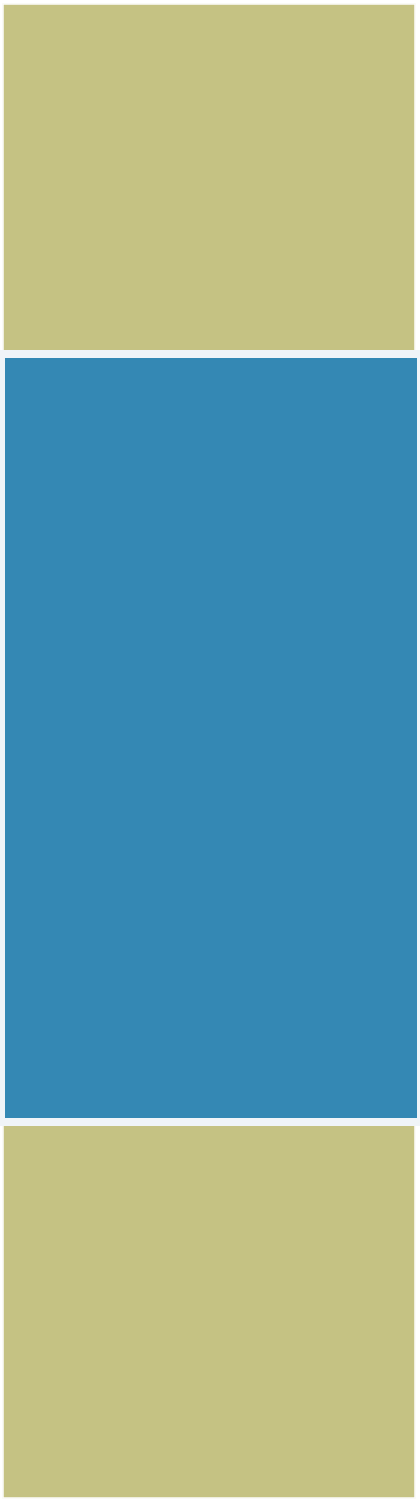

Palo Alto Networks

# PAN-OS® Administrator's Guide
Version 7.0

## Contact Information

**Corporate Headquarters:**

Palo Alto Networks

4401 Great America Parkway

Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-us

## About this Guide

This guide takes you through the configuration and maintenance of your Palo Alto Networks next-generation firewall. For additional information, refer to the following resources:

- For information on the additional capabilities and for instructions on configuring the features on the firewall, refer to https://www.paloaltonetworks.com/documentation.
- For access to the knowledge base, complete documentation set, discussion forums, and videos, refer to https://live.paloaltonetworks.com.
- For contacting support, for information on support programs, to manage your account or devices, or to open a support case, refer to https://www.paloaltonetworks.com/support/tabs/overview.html.
- For the most current PAN-OS 7.0 release notes, see https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os-release-notes.html.

To provide feedback on the documentation, please write to us at: documentation@paloaltonetworks.com.

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2015–2016 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at http://www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

**Revision Date: March 23, 2016**

# Table of Contents


## Getting Started . . . . . 15

Integrate the Firewall into Your Management Network . . . . . 16
- Determine Your Management Strategy . . . . . 17
- Perform Initial Configuration . . . . . 18
- Set Up Network Access for External Services . . . . . 22
- Register the Firewall . . . . . 26
- Activate Licenses and Subscriptions . . . . . 27
- Manage Content Updates . . . . . 29
- Install Software Updates . . . . . 32

Create the Security Perimeter . . . . . 33
- Basic Interface Deployments . . . . . 34
- About Security Policy . . . . . 36
- Plan the Deployment . . . . . 39
- Configure Interfaces and Zones . . . . . 40
- Set Up Basic Security Policies . . . . . 42

Enable Basic Threat Prevention Features . . . . . 48
- Enable WildFire . . . . . 49
- Scan Traffic for Threats . . . . . 51
- Control Access to Web Content . . . . . 56

Best Practices for Completing the Firewall Deployment . . . . . 59

## Device Management . . . . . 61

Management Interfaces . . . . . 62
- Use the Web Interface . . . . . 63
- Use the Command Line Interface (CLI) . . . . . 70
- Use the XML API . . . . . 72

Manage Firewall Administrators . . . . . 74
- Administrative Roles . . . . . 75
- Administrative Authentication . . . . . 77
- Configure Administrative Accounts and Authentication . . . . . 78
- Configure an Administrative Account . . . . . 79
- Configure Kerberos SSO and External or Local Authentication for Administrators . . . . . 80
- Configure Certificate-Based Administrator Authentication to the Web Interface . . . . . 82
- Configure SSH Key-Based Administrator Authentication to the CLI . . . . . 83
- Configure RADIUS Vendor-Specific Attributes for Administrator Authentication . . . . . 84

Reference: Web Interface Administrator Access . . . . . 86
- Web Interface Access Privileges . . . . . 86
- Panorama Web Interface Access . . . . . 126

Reference: Port Numbers Used by Palo Alto Networks Devices . . . . . 129
- Ports Used for Management Functions . . . . . 129
- Ports Used for HA . . . . . 130
- Ports Used for Panorama . . . . . 131
- Ports Used for User-ID . . . . . 132

Reset the Firewall to Factory Default Settings . . . . . 134


© Palo Alto Networks, Inc.

## Authentication . . . . . 135

Configure Kerberos Single Sign-On . . . . . 136
Configure External Authentication . . . . . 137
- Configure Authentication Server Profiles . . . . . 138
- Configure a RADIUS Server Profile . . . . . 139
- RADIUS Vendor-Specific Attributes for Palo Alto Networks Devices . . . . . 140
- Configure a TACACS+ Server Profile . . . . . 141
- Configure an LDAP Server Profile . . . . . 142
- Configure a Kerberos Server Profile . . . . . 143
- Set CHAP and PAP Authentication for RADIUS and TACACS+ Servers . . . . . 144
- Configure an Authentication Profile and Sequence . . . . . 145
- Enable External Authentication for Users and Services . . . . . 148

Test Authentication Server Connectivity . . . . . 149
- Run the Test Authentication Command . . . . . 150
- Local Database Authentication Profile Use Case . . . . . 151
- RADIUS Authentication Profile Use Case . . . . . 152
- TACACS+ Authentication Profile Use Case . . . . . 154
- LDAP Authentication Profile Use Case . . . . . 157
- Kerberos Authentication Profile Use Case . . . . . 159

Troubleshoot Authentication Issues . . . . . 161

## Certificate Management . . . . . 163

Keys and Certificates . . . . . 164
Certificate Revocation . . . . . 166
- Certificate Revocation List (CRL) . . . . . 166
- Online Certificate Status Protocol (OCSP) . . . . . 167

Certificate Deployment . . . . . 168
Set Up Verification for Certificate Revocation Status . . . . . 169
- Configure an OCSP Responder . . . . . 169
- Configure Revocation Status Verification of Certificates Used for User/Device Authentication . . . 170
- Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption . . . . . 170

Configure the Master Key . . . . . 172
Obtain Certificates . . . . . 173
- Create a Self-Signed Root CA Certificate . . . . . 174
- Generate a Certificate on the Device . . . . . 175
- Import a Certificate and Private Key . . . . . 177
- Obtain a Certificate from an External CA . . . . . 178

Export a Certificate and Private Key . . . . . 180
Configure a Certificate Profile . . . . . 181
Configure an SSL/TLS Service Profile . . . . . 183
Configure the Key Size for SSL Forward Proxy Server Certificates . . . . . 184
Revoke and Renew Certificates . . . . . 185
- Revoke a Certificate . . . . . 185
- Renew a Certificate . . . . . 185



© Palo Alto Networks, Inc.

Secure Keys with a Hardware Security Module . . . 186
Set up Connectivity with an HSM . . . 187
Encrypt a Master Key Using an HSM . . . 193
Store Private Keys on an HSM. . . . 195
Manage the HSM Deployment. . . . 196

High Availability . . . 197

HA Overview . . . 198

HA Concepts . . . 199
HA Modes . . . 199
HA Links and Backup Links . . . 199
Device Priority and Preemption. . . . 202
Failover Triggers . . . 202
HA Timers. . . . 203

Set Up Active/Passive HA. . . . 206
Prerequisites for Active/Passive HA . . . 207
Configuration Guidelines for Active/Passive HA . . . 208
Configure Active/Passive HA . . . 210
Define HA Failover Conditions . . . 216
Verify Failover. . . . 218

Reference: HA Synchronization. . . . 219
What Settings Don't Sync in Active/Passive HA? . . . 219
What Settings Don't Sync in Active/Active HA? . . . 221
Synchronization of System Runtime Information. . . . 223

HA Resources . . . 225

Monitoring . . . 227

Use the Dashboard . . . 228

Use the Application Command Center . . . 229
ACC—First Look . . . 230
ACC Tabs . . . 232
ACC Widgets. . . . 233
Widget Descriptions . . . 235
ACC Filters . . . 239
Interact with the ACC. . . . 241
Use Case: ACC—Path of Information Discovery . . . 244

App Scope. . . . 250
Summary Report . . . 251
Change Monitor Report . . . 252
Threat Monitor Report . . . 254
Threat Map Report . . . 255
Network Monitor Report . . . 256
Traffic Map Report . . . 257

Use the Automated Correlation Engine . . . 258
Automated Correlation Engine Concepts . . . 259
View the Correlated Objects . . . 260
Interpret Correlated Events . . . 261
Use the Compromised Hosts Widget in the ACC . . . 264

© Palo Alto Networks, Inc.

Take Packet Captures. . . . . 265
Disable Hardware Offload. . . . . 266
Take a Custom Packet Capture . . . . . 266
Take a Threat Packet Capture . . . . . 271
Take an Application Packet Capture . . . . . 272
Monitor Applications and Threats . . . . . 278
Monitor and Manage Logs. . . . . 279
View the Log Files . . . . . 280
Filter Log Data . . . . . 283
Configure Log Storage Quotas and Expiration Periods. . . . . 284
Log Severity Levels and WildFire Verdicts . . . . . 285
Schedule Log Exports to an SCP or FTP Server . . . . . 287
Manage Reporting . . . . . 288
Report Types . . . . . 289
View Reports . . . . . 290
Configure the Report Expiration Period. . . . . 291
Disable Predefined Reports . . . . . 292
Generate Custom Reports . . . . . 293
Generate Botnet Reports . . . . . 299
Manage PDF Summary Reports . . . . . 301
Generate User/Group Activity Reports . . . . . 303
Manage Report Groups . . . . . 305
Schedule Reports for Email Delivery . . . . . 306
Use External Services for Monitoring. . . . . 307
Configure Log Forwarding . . . . . 308
Configure Email Alerts . . . . . 311
Use Syslog for Monitoring. . . . . 312
Configure Syslog Monitoring. . . . . 313
Syslog Field Descriptions. . . . . 316
SNMP Monitoring and Traps . . . . . 333
SNMP for Palo Alto Networks Devices . . . . . 334
Use an SNMP Manager to Explore MIBs and Objects . . . . . 336
Enable SNMP Services for Firewall-Secured Network Elements . . . . . 340
Monitor Device Statistics Using SNMP . . . . . 341
Forward Traps to an SNMP Manager. . . . . 344
Supported MIBs . . . . . 346
NetFlow Monitoring . . . . . 354
Configure NetFlow Exports . . . . . 354
NetFlow Templates . . . . . 355
Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors . . . . . 359

User-ID . . . . . 361

User-ID Overview . . . . . 362
User-ID Concepts . . . . . 364
Group Mapping . . . . . 364
User Mapping . . . . . 365
Enable User-ID . . . . . 369


© Palo Alto Networks, Inc.

Map Users to Groups . . . 370
Map IP Addresses to Users . . . 373
Configure User Mapping Using the Windows User-ID Agent . . . 374
Configure User Mapping Using the PAN-OS Integrated User-ID Agent . . . 380
Configure User-ID to Receive User Mappings from a Syslog Sender . . . 383
Map IP Addresses to Usernames Using Captive Portal . . . 393
Configure User Mapping for Terminal Server Users . . . 401
Send User Mappings to User-ID Using the XML API . . . 409
Configure a Firewall to Share User Mapping Data with Other Firewalls . . . 410
Enable User- and Group-Based Policy . . . 412
Enable Policy for Users with Multiple Accounts . . . 414
Verify the User-ID Configuration . . . 416
Deploy User-ID in a Large-Scale Network . . . 419
Windows Log Forwarding and Global Catalog Servers . . . 419
Plan Your User-ID Implementation for a Large-Scale Network . . . 420
Configure Windows Log Forwarding . . . 421
Configure User-ID for a Large-Scale Network . . . 421

App-ID . . . 425
App-ID Overview . . . 426
Manage Custom or Unknown Applications . . . 427
Manage New App-IDs Introduced in Content Releases . . . 429
Review New App-IDs . . . 430
Review New App-IDs Since Last Content Version . . . 431
Review New App-ID Impact on Existing Policy Rules . . . 432
Disable or Enable App-IDs . . . 433
Prepare Policy Updates For Pending App-IDs . . . 434
Use Application Objects in Policy . . . 436
Create an Application Group . . . 436
Create an Application Filter . . . 437
Create a Custom Application . . . 438
Applications with Implicit Support . . . 443
Application Level Gateways . . . 446
Disable the SIP Application-level Gateway (ALG) . . . 447

Threat Prevention . . . 449
Set Up Security Profiles and Policies . . . 450
Set Up Antivirus, Anti-Spyware, and Vulnerability Protection . . . 451
Set Up Data Filtering . . . 453
Set Up File Blocking . . . 457
Prevent Brute Force Attacks . . . 459
Customize the Action and Trigger Conditions for a Brute Force Signature . . . 460
Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions . . . 463
Enable Passive DNS Collection for Improved Threat Intelligence . . . 466

© Palo Alto Networks, Inc.

Use DNS Queries to Identify Infected Hosts on the Network . . . 467
DNS Sinkholing . . . 468
Configure DNS Sinkholing . . . 469
Identify Infected Hosts . . . 473
Content Delivery Network Infrastructure for Dynamic Updates . . . 475
Threat Prevention Resources . . . 477

Decryption . . . 479
Decryption Overview . . . 480
Decryption Concepts . . . 481
Keys and Certificates for Decryption Policies . . . 482
SSL Forward Proxy . . . 484
SSL Inbound Inspection . . . 486
SSH Proxy . . . 487
Decryption Exceptions . . . 488
Decryption Mirroring . . . 489
Configure SSL Forward Proxy . . . 490
Configure SSL Inbound Inspection . . . 495
Configure SSH Proxy . . . 497
Configure Decryption Exceptions . . . 498
Exclude Traffic from Decryption . . . 498
Exclude a Server from Decryption . . . 499
Enable Users to Opt Out of SSL Decryption . . . 500
Configure Decryption Port Mirroring . . . 502
Temporarily Disable SSL Decryption . . . 504

URL Filtering . . . 505
URL Filtering Overview . . . 506
URL Filtering Vendors . . . 507
Interaction Between App-ID and URL Categories . . . 508
PAN-DB Private Cloud . . . 509
URL Filtering Concepts . . . 511
URL Categories . . . 512
URL Filtering Profile . . . 513
URL Filtering Profile Actions . . . 514
Block and Allow Lists . . . 515
Safe Search Enforcement . . . 516
Container Pages . . . 518
HTTP Header Logging . . . 519
URL Filtering Response Pages . . . 520
URL Category as Policy Match Criteria . . . 522
PAN-DB Categorization . . . 523
PAN-DB URL Categorization Components . . . 523
PAN-DB URL Categorization Workflow . . . 525


© Palo Alto Networks, Inc.

Enable a URL Filtering Vendor . . . 526
Enable PAN-DB URL Filtering . . . 527
Enable BrightCloud URL Filtering . . . 528
Determine URL Filtering Policy Requirements . . . 530
Monitor Web Activity . . . 532
Monitor Web Activity of Network Users . . . 533
View the User Activity Report . . . 536
Configure Custom URL Filtering Reports . . . 538
Configure URL Filtering . . . 539
Customize the URL Filtering Response Pages . . . 541
Configure URL Admin Override . . . 542
Enable Safe Search Enforcement . . . 544
Block Search Results that are not Using Strict Safe Search Settings . . . 544
Enable Transparent Safe Search Enforcement . . . 548
Set Up the PAN-DB Private Cloud . . . 552
URL Filtering Use Case Examples . . . 557
Use Case: Control Web Access . . . 558
Use Case: Use URL Categories for Policy Matching . . . 562
Troubleshoot URL Filtering . . . 564
Problems Activating PAN-DB . . . 564
PAN-DB Cloud Connectivity Issues . . . 565
URLs Classified as Not-Resolved . . . 566
Incorrect Categorization . . . 567
URL Database Out of Date . . . 568

Quality of Service . . . 569
QoS Overview . . . 570
QoS Concepts . . . 572
QoS for Applications and Users . . . 572
QoS Profile . . . 573
QoS Classes . . . 574
QoS Policy . . . 575
QoS Egress Interface . . . 576
QoS Clear Text and Tunneled Traffic . . . 577
Configure QoS . . . 578
Configure QoS for a Virtual System . . . 583
Enforce QoS Based on DSCP Classification . . . 590
QoS Use Cases . . . 593
Use Case: QoS for a Single User . . . 594
Use Case: QoS for Voice and Video Applications . . . 597

VPNs . . . 601
VPN Deployments . . . 602
Site-to-Site VPN Overview . . . 603
Site-to-Site VPN Concepts . . . 604

© Palo Alto Networks, Inc.

IKE Gateway . . . 604
Tunnel Interface . . . 604
Tunnel Monitoring . . . 605
Internet Key Exchange (IKE) for VPN . . . 605
IKEv2 . . . 608
Set Up Site-to-Site VPN . . . 612
Set Up an IKE Gateway . . . 613
Define Cryptographic Profiles . . . 620
Set Up an IPSec Tunnel . . . 623
Set Up Tunnel Monitoring . . . 627
Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel . . . 628
Test VPN Connectivity . . . 630
Interpret VPN Error Messages . . . 631
Site-to-Site VPN Quick Configs . . . 632
Site-to-Site VPN with Static Routing . . . 633
Site-to-Site VPN with OSPF . . . 638
Site-to-Site VPN with Static and Dynamic Routing . . . 644

Large Scale VPN (LSVPN) . . . 651

LSVPN Overview . . . 652
Create Interfaces and Zones for the LSVPN . . . 653
Enable SSL Between GlobalProtect LSVPN Components . . . 655
About Certificate Deployment . . . 655
Deploy Server Certificates to the GlobalProtect LSVPN Components . . . 655
Configure the Portal to Authenticate Satellites . . . 659
Configure GlobalProtect Gateways for LSVPN . . . 661
Prerequisite Tasks . . . 661
Configure the Gateway . . . 661
Configure the GlobalProtect Portal for LSVPN . . . 664
Prerequisite Tasks . . . 664
Configure the Portal . . . 664
Define the Satellite Configurations . . . 665
Prepare the Satellite Device to Join the LSVPN . . . 668
Verify the LSVPN Configuration . . . 671
LSVPN Quick Configs . . . 672
Basic LSVPN Configuration with Static Routing . . . 673
Advanced LSVPN Configuration with Dynamic Routing . . . 677

Networking . . . 681

Interface Deployments . . . 682
Virtual Wire Deployments . . . 682
Layer 2 Deployments . . . 685
Layer 3 Deployments . . . 685
Tap Mode Deployments . . . 686
Virtual Routers . . . 687


© Palo Alto Networks, Inc.

Static Routes . . . 689
RIP . . . 691
OSPF . . . 693
OSPF Concepts . . . 693
Configure OSPF . . . 696
Configure OSPFv3 . . . 701
Configure OSPF Graceful Restart . . . 704
Confirm OSPF Operation . . . 705
BGP . . . 708
Session Settings and Timeouts . . . 713
Transport Layer Sessions . . . 714
TCP . . . 715
UDP . . . 719
ICMP . . . 720
Configure Session Timeouts . . . 721
Configure Session Settings . . . 723
Prevent TCP Split Handshake Session Establishment . . . 725
DHCP . . . 726
DHCP Overview . . . 727
DHCP Messages . . . 728
DHCP Addressing . . . 729
DHCP Options . . . 731
Firewall as a DHCP Server and Client . . . 734
Configure an Interface as a DHCP Server . . . 735
Configure an Interface as a DHCP Client . . . 740
Configure an Interface as a DHCP Relay Agent . . . 742
Monitor and Troubleshoot DHCP . . . 743
NAT . . . 745
NAT Policy Rules . . . 746
Source NAT and Destination NAT . . . 749
NAT Rule Capacities . . . 751
Dynamic IP and Port NAT Oversubscription . . . 752
Dataplane NAT Memory Statistics . . . 754
Configure NAT . . . 755
NAT Configuration Examples . . . 765
NPTv6 . . . 773
NPTv6 Overview . . . 774
How NPTv6 Works . . . 776
NDP Proxy . . . 778
NPTv6 and NDP Proxy Example . . . 780
Create an NPTv6 Policy . . . 782
LACP . . . 784
LACP Settings . . . 784
Configure LACP . . . 787
ECMP . . . 790
ECMP Load-Balancing Algorithms . . . 791
ECMP Platform, Interface, and IP Routing Support . . . 793
HA Active/Active Failover Behavior with ECMP . . . 794

© Palo Alto Networks, Inc.

Configure ECMP on a Virtual Router . . . 795
Enable ECMP for Multiple BGP Autonomous Systems . . . 797
Verify ECMP . . . 799
LLDP . . . 800
LLDP Overview . . . 801
Supported TLVs in LLDP . . . 802
LLDP Syslog Messages and SNMP Traps . . . 804
Configure LLDP . . . 805
View LLDP Settings and Status . . . 807
Clear LLDP Statistics . . . 810

Policy . . . 811

Policy Types . . . 812
Security Policy . . . 813
Components of a Security Policy Rule . . . 814
Security Policy Best Practices . . . 817
Policy Objects . . . 818
Security Profiles . . . 819
Antivirus Profiles . . . 820
Anti-Spyware Profiles . . . 821
Vulnerability Protection Profiles . . . 822
URL Filtering Profiles . . . 823
Data Filtering Profiles . . . 824
File Blocking Profiles . . . 826
WildFire Analysis Profiles . . . 827
DoS Protection Profiles . . . 828
Zone Protection Profiles . . . 829
Security Profile Group . . . 830
Enumeration of Rules Within a Rulebase . . . 834
Move or Clone a Policy Rule or Object to a Different Virtual System . . . 836
Use Tags to Group and Visually Distinguish Objects . . . 837
Create and Apply Tags . . . 838
Modify Tags . . . 839
Use the Tag Browser . . . 840
Use a Dynamic Block List in Policy . . . 844
View the IP Address Limit For Your Firewall Model . . . 844
Formatting Guidelines for Dynamic Block Lists . . . 845
Enforce Policy with a Dynamic Block List . . . 845
View the List of IP addresses in the Dynamic Block List . . . 846
Retrieve a Dynamic Block List from Web Server . . . 847
Register IP Addresses and Tags Dynamically . . . 848
Monitor Changes in the Virtual Environment . . . 849
Enable VM Monitoring to Track Changes on the Virtual Network . . . 850
Attributes Monitored in the AWS and VMware Environments . . . 853
Use Dynamic Address Groups in Policy . . . 854
CLI Commands for Dynamic IP Addresses and Tags . . . 857


© Palo Alto Networks, Inc.

Identify Users Connected through a Proxy Server . . . 859
Use XFF Values for Policies and Logging Source Users . . . 859
Add XFF Values to URL Filtering Logs . . . 860
Policy-Based Forwarding . . . 861
PBF . . . 862
Create a Policy-Based Forwarding Rule . . . 864
Use Case: PBF for Outbound Access with Dual ISPs . . . 866
DoS Protection Against Flooding of New Sessions . . . 874
DoS Protection Against Flooding of New Sessions . . . 875
Configure DoS Protection Against Flooding of New Sessions . . . 879
Use the CLI to End a Single Attacking Session . . . 882
Identify Sessions That Use an Excessive Percentage of the Packet Buffer . . . 883
Discard a Session Without a Commit . . . 885

Virtual Systems . . . 887

Virtual Systems Overview . . . 888
Virtual System Components and Segmentation . . . 888
Benefits of Virtual Systems . . . 889
Use Cases for Virtual Systems . . . 889
Platform Support and Licensing for Virtual Systems . . . 890
Administrative Roles for Virtual Systems . . . 890
Shared Objects for Virtual Systems . . . 890
Communication Between Virtual Systems . . . 891
Inter-VSYS Traffic That Must Leave the Firewall . . . 892
Inter-VSYS Traffic That Remains Within the Firewall . . . 893
Inter-VSYS Communication Uses Two Sessions . . . 896
Shared Gateway . . . 897
External Zones and Shared Gateway . . . 897
Networking Considerations for a Shared Gateway . . . 898
Service Routes for Virtual Systems . . . 899
Use Cases for Service Routes for a Virtual System . . . 900
PA-7000 Series Firewall LPC Support for Per-Virtual System Paths to Logging Servers . . . 901
DNS Proxy Object . . . 902
DNS Server Profile . . . 903
Multi-Tenant DNS Deployments . . . 904
Configure Virtual Systems . . . 905
Configure Inter-Virtual System Communication within the Firewall . . . 908
Configure a Shared Gateway . . . 909
Customize Service Routes for a Virtual System . . . 910
Customize Service Routes to Services for Virtual Systems . . . 910
Configure a PA-7000 Series Firewall for Logging Per Virtual System . . . 912
Configure a DNS Proxy Object . . . 913
Configure a DNS Server Profile . . . 914
Configure Administrative Access Per Virtual System or Device . . . 916

© Palo Alto Networks, Inc.

DNS Resolution—Three Use Cases . . . . . 918
Use Case 1: Firewall Requires DNS Resolution for Management Purposes . . . . . 918
Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System . . . . . 921
Use Case 3: Firewall Acts as DNS Proxy Between Client and Server . . . . . 925
Virtual System Functionality with Other Features . . . . . 927

Certifications . . . . . 929
Enable FIPS and Common Criteria Support . . . . . 930
CCEAL4 Security Functions . . . . . 931


© Palo Alto Networks, Inc.

# Getting Started

The following topics provide detailed steps to help you deploy a new Palo Alto Networks next-generation firewall. They provide details for integrating a new firewall into your network and configuring basic security policies and threat prevention features.

After you perform the basic configuration steps required to integrate the firewall into your network, you can use the rest of the topics in this guide to help you deploy the comprehensive enterprise security platform features as necessary to address your network security needs.

- Integrate the Firewall into Your Management Network
- Create the Security Perimeter
- Enable Basic Threat Prevention Features
- Best Practices for Completing the Firewall Deployment

# Integrate the Firewall into Your Management Network

All Palo Alto Networks firewalls provide an out-of-band management port (MGT) that you can use to perform the firewall administration functions. By using the MGT port, you separate the management functions of the firewall from the data processing functions, safeguarding access to the firewall and enhancing performance. When using the web interface, you must perform all initial configuration tasks from the MGT port even if you plan to use an in-band port for managing your device going forward.

Some management tasks, such as retrieving licenses and updating the threat and application signatures on the firewall require access to the Internet. If you do not want to enable external access to your MGT port, you will need to either set up a data port to provide access to required external services or plan to manually upload updates regularly.

The following topics describe how to perform the initial configuration steps that are necessary to integrate a new firewall into the management network and deploy it in a basic security configuration.

- Determine Your Management Strategy
- Perform Initial Configuration
- Set Up Network Access for External Services
- Register the Firewall
- Activate Licenses and Subscriptions
- Manage Content Updates
- Install Software Updates

The following topics describe how to integrate a single Palo Alto Networks next-generation firewall into your network. However, for redundancy, consider deploying a pair of firewalls in a High Availability configuration.

 © Palo Alto Networks, Inc.

## Determine Your Management Strategy

The Palo Alto Networks firewall can be configured and managed locally or it can be managed centrally using Panorama, the Palo Alto Networks centralized security management system. If you have six or more firewalls deployed in your network, use Panorama to achieve the following benefits:

- Reduce the complexity and administrative overhead in managing configuration, policies, software and dynamic content updates. Using device groups and templates on Panorama, you can effectively manage device specific configuration locally on a device and enforce shared policies across all devices or device groups.
- Aggregate data from all managed firewalls and gain visibility across all the traffic on your network. The Application Command Center (ACC) on Panorama provides a single glass pane for unified reporting across all the firewalls, allowing you to centrally analyze, investigate and report on network traffic, security incidents and administrative modifications.

The procedures that follow describe how to manage the firewall using the local web interface. If you want to use Panorama for centralized management, first Perform Initial Configuration and verify that the firewall can establish a connection to Panorama. From that point on you can use Panorama to configure your firewall centrally.

## Perform Initial Configuration

By default, the firewall has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other firewall configuration tasks. You must perform these initial configuration tasks either from the MGT interface, even if you do not plan to use this interface for your firewall management, or using a direct serial connection to the console port on the device.

| **Set Up Network Access to the Firewall** | | |
|---|---|---|
| Step 1 | Gather the required information from your network administrator. | • IP address for MGT port<br>• Netmask<br>• Default gateway<br>• DNS server address |
| Step 2 | Connect your computer to the firewall. | You can connect to the firewall in one of the following ways:<br>• Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1). Wait a few minutes for the boot-up sequence to complete; when the device is ready, the prompt changes to the name of the firewall, for example `PA-500 login`.<br>• Connect an RJ-45 Ethernet cable from your computer to the MGT port on the firewall. From a browser, go to `https://192.168.1.1`. Note that you may need to change the IP address on your computer to an address in the 192.168.1.0 network, such as 192.168.1.2, in order to access this URL. |
| Step 3 | When prompted, log in to the firewall. | You must log in using the default username and password (admin/admin). The firewall will begin to initialize. |
| Step 4 | Configure the MGT interface. | 1. Select **Device > Setup > Management** and then edit the Management Interface Settings.<br>2. Enter the **IP Address**, **Netmask**, and **Default Gateway**.<br> To prevent unauthorized access to the management interface, it is a best practice to **Add** the **Permitted IP Addresses** from which an administrator can access the MGT interface.<br>3. Set the **Speed** to **auto-negotiate**.<br>4. Select which management services to allow on the interface.<br> Make sure **Telnet** and **HTTP** are not selected because these services use plaintext and are not as secure as the other services and could compromise administrator credentials.<br>5. Click **OK**. |

 © Palo Alto Networks, Inc.

| Set Up Network Access to the Firewall (Continued) | | |
|---|---|---|
| Step 5 | Configure general firewall settings as needed.<br> <br>As a best practice, add a login banner that indicates that access to the firewall is restricted to prevent unauthorized users from accessing the management functions. It is a good idea to avoid using welcoming verbiage and to run your messaging by your legal department to ensure that you are providing adequate warning that unauthorized access is prohibited. | 1. Select **Device > Setup > Management** and edit the General Settings.<br>2. Enter a **Hostname** for the firewall and enter your network **Domain** name. The domain name is just a label; it will not be used to join the domain.<br>3. Enter **Login Banner** text that informs users who are attempting to log in that they are that they must have authorization to access the firewall management functions.<br>4. Enter the **Latitude** and **Longitude** to enable accurate placement of the firewall on the world map.<br>5. Click **OK**. |
| Step 6 | Configure DNS, update server, and proxy server settings.<br> <br>You must manually configure at least one DNS server on the firewall or it will not be able to resolve hostnames; it will not use DNS server settings from another source, such as an ISP. | 1. Select **Device > Setup > Services**.<br>• For multi-virtual system platforms, select **Global** and edit the Services section.<br>• For single virtual system platforms, edit the Services section.<br>2. On the **Services** tab, for **DNS**, click one of the following:<br>• **Servers**—Enter the **Primary DNS Server** address and **Secondary DNS Server** address.<br>• **DNS Proxy Object**—From the drop-down, select the **DNS Proxy** that you want to use to configure global DNS services, or click **DNS Proxy** to configure a new DNS proxy object.<br>3. (Optional) For **Update Server**, enter the IP address or host name of the server from which to download updates from Palo Alto Networks. The current value is updates.paloaltonetworks.com. Do not change the Update Server unless instructed by Technical Support.<br>4. (Optional) Click **Verify Update Server Identity** for an extra level of validation. The firewall will check the update server's SSL certificate to ensure that it was signed by a trusted authority.<br>5. (Optional) If the firewall needs to use a proxy server to reach Palo Alto Networks update services, in the **Proxy Server** window, enter:<br>• **Server**—IP address or host name of the proxy server.<br>• **Port**—Port for the proxy server. Range: 1-65535.<br>• **User**—Username to access the server.<br>• **Password**—Password for the user to access the proxy server. Re-enter the password at **Confirm Password**.<br>6. Click **OK**. |

© Palo Alto Networks, Inc.

| Set Up Network Access to the Firewall (Continued) | |
|---|---|
| Step 7 Configure date and time (NTP) settings. | 1. Select **Device > Setup > Services**.<br>• For multi-virtual system platforms, select **Global** and edit the Services section.<br>• For single virtual system platforms, edit the Services section.<br>2. On the **NTP** tab, to use the virtual cluster of time servers on the Internet, enter the hostname **pool.ntp.org** as the **Primary NTP Server** or enter the IP address of your primary NTP server.<br>3. (Optional) Enter a **Secondary NTP Server** address.<br>4. (Optional) To authenticate time updates from the NTP server(s), for **Authentication Type**, select one of the following for each server:<br>• **None**—(Default) Disables NTP authentication.<br>• **Symmetric Key**—Firewall uses symmetric key exchange (shared secrets) to authenticate time updates.<br>– **Key ID**—Enter the Key ID (1-65534).<br>– **Algorithm**—Select the algorithm to use in NTP authentication (**MD5** or **SHA1**).<br>• **Autokey**—Firewall uses autokey (public key cryptography) to authenticate time updates.<br>5. Click **OK**. |
| Step 8 Set a secure password for the admin account. | 1. Select **Device > Administrators.**<br>2. Select the **admin** role.<br>3. Enter the current default password and the new password.<br>4. Click **OK** to save your settings. |
| Step 9 Commit your changes.<br>When the configuration changes are saved, you will lose connectivity to the web interface because the IP address will have changed. | Click **Commit**. The device may take up to 90 seconds to save your changes.<br> |
| Step 10 Connect the firewall to your network. | 1. Disconnect the firewall from your computer.<br>2. Connect the MGT port to a switch port on your management network using an RJ-45 Ethernet cable. Make sure that the switch port you cable the firewall to is configured for auto-negotiation. |
| Step 11 Open an SSH management session to the firewall. | Using a terminal emulation software, such as PuTTY, launch an SSH session to the firewall using the new IP address you assigned to it. |

 © Palo Alto Networks, Inc.

| Set Up Network Access to the Firewall (Continued) | |
|---|---|
| Step 12 Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server, in one of the following ways:<br>• If you do not want to allow external network access to the MGT interface, you will need to set up a data port to retrieve required service updates. Continue to Set Up Network Access for External Services.<br>• If you do plan to allow external network access to the MGT interface, verify that you have connectivity and then proceed to Register the Firewall and Activate Licenses and Subscriptions. | If you cabled your MGT port for external network access, verify that you have access to and from the firewall by using the ping utility from the CLI. Make sure you have connectivity to the default gateway, DNS server, and the Palo Alto Networks Update Server as shown in the following example:<br>`admin@PA-200> ping host updates.paloaltonetworks.com`<br>`PING updates.paloaltonetworks.com (67.192.236.252) 56(84) bytes of data.`<br>`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=40.5 ms`<br>`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=53.6 ms`<br>`64 bytes from 67.192.236.252 : icmp_seq=1 ttl=243 time=79.5 ms`<br>After you have verified connectivity, press Ctrl+C to stop the pings. |

## Set Up Network Access for External Services

By default, the firewall uses the MGT interface to access remote services, such as DNS servers, content updates, and license retrieval. If you do not want to enable external network access to your management network, you must set up a data port to provide access to these required external services.

This task requires familiarity with firewall interfaces, zones, and policies. For more information on these topics, see Create the Security Perimeter.

For information on setting up network access to external services on a virtual system basis rather than a global basis, see Per-Virtual System Service Routes.

| Set Up a Data Port for Access to External Services | | |
|---|---|---|
| Step 1 | Decide which port you want to use for access to external services and connect it to your switch or router port. | The interface you use must have a static IP address. |
| Step 2 | Log in to the web interface. | Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page. |
| Step 3 | (Optional) The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and zones). If you do not plan to use this virtual wire configuration, you must manually delete the configuration to prevent it from interfering with other interface settings you define. | You must delete the configuration in the following order:<br>1. To delete the default security policy, select **Policies > Security**, select the rule, and click **Delete**.<br>2. Next, delete the default virtual wire by selecting **Network > Virtual Wires**, selecting the virtual wire and clicking **Delete**.<br>3. To delete the default trust and untrust zones, select **Network > Zones**, select each zone and click **Delete**.<br>4. Finally, delete the interface configurations by selecting **Network > Interfaces** and then select each interface (ethernet1/1 and ethernet1/2) and click **Delete**.<br>5. **Commit** the changes. |

 © Palo Alto Networks, Inc.

| Set Up a Data Port for Access to External Services (Continued) | |
|---|---|
| Step 4 Configure the interface. | 1. Select **Network > Interfaces** and select the interface that corresponds to the port you cabled in Step 1.<br>2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**.<br>4. In the Zone dialog, define a **Name** for new zone, for example L3-trust, and then click **OK**.<br>5. Select the **IPv4** tab, select the **Static** radio button, and click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.254/24.<br>6. Select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**.<br>7. Enter a **Name** for the profile, such as allow_ping, and then select the services you want to allow on the interface. For the purposes of allowing access to the external services, you probably only need to enable **Ping** and then click **OK**.<br> These services provide management access to the device, so only select the services that correspond to the management activities you want to allow on this interface. For example, if you plan to use the MGT interface for device configuration tasks through the web interface or CLI, you would not want to enable HTTP, HTTPS, SSH, or Telnet so that you could prevent unauthorized access through this interface (and if you did allow those services, you should limit access to a specific set of **Permitted IP Addresses**).<br>8. To save the interface configuration, click **OK**. |

**Set Up a Data Port for Access to External Services (Continued)**

Step 5 Because the firewall uses the MGT interface by default to access the external services it requires, you must change the interface the firewall uses to send these requests by editing the service routes.

1. Select **Device > Setup > Services > Service Route Configuration**.

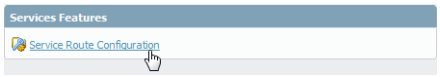


For the purposes of activating your licenses and getting the most recent content and software updates, you will want to change the service route for **DNS**, **Palo Alto Updates**, **URL Updates**, and **WildFire**.

2. Click the **Customize** radio button, and select one of the following:
   - For a predefined service, select **IPv4** or **IPv6** and click the link for the service for which you want to modify the **Source Interface** and select the interface you just configured.

     If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.
   - To create a service route for a custom destination, select **Destination**, and click **Add**. Enter a **Destination** name and select a **Source Interface**. If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.
3. Click **OK** to save the settings.
4. Repeat steps 2-3 above for each service route you want to modify.
5. **Commit** your changes.

 © Palo Alto Networks, Inc.

| Set Up a Data Port for Access to External Services (Continued) | |
|---|---|
| Step 6 Configure an external-facing interface and an associated zone and then create security and NAT policy rules to allow the firewall to send service requests from the internal zone to the external zone. | 1. Select **Network > Interfaces** and then select your external-facing interface. Select **Layer3** as the **Interface Type**, **Add** the **IP** address (on the **IPv4** or **IPv6** tab), and create the associated **Security Zone** (on the **Config** tab), such as l3-untrust. You do not need to set up management services on this interface.<br>2. To set up a security rule that allows traffic from your internal network to the Palo Alto Networks update server, select **Policies > Security** and click **Add**. For the purposes of initial configuration, you can create a simple rule that allows all traffic from l3-trust to l3-untrust as follows:<br> <br>3. If you are using a private IP address on the internal-facing interface, you will need to create a source NAT rule to translate the address to a publicly routable address. Select **Policies > NAT** and then click **Add**. At a minimum you must define a name for the rule (**General** tab), specify a source and destination zone, l3-trust to l3-untrust in this case (**Original Packet** tab), and define the source address translation settings (**Translated Packet** tab) and then click **OK**.<br> <br>4. **Commit** your changes. |
| Step 7 Verify that you have connectivity from the data port to the external services, including the default gateway, and the Palo Alto Networks Update Server.<br><br>After you verify you have the required network connectivity, continue to Register the Firewall and Activate Licenses and Subscriptions. | Launch the CLI and use the ping utility to verify that you have connectivity. Keep in mind that by default pings are sent from the MGT interface, so in this case you must specify the source interface for the ping requests as follows:<br>`admin@PA-200> ping source 192.168.1.254 host updates.paloaltonetworks.com`<br>`PING updates.paloaltonetworks.com (67.192.236.252) from 192.168.1.254 : 56(84) bytes of data.`<br>`64 bytes from 67.192.236.252: icmp_seq=1 ttl=242 time=56.7 ms`<br>`64 bytes from 67.192.236.252: icmp_seq=2 ttl=242 time=47.7 ms`<br>`64 bytes from 67.192.236.252: icmp_seq=3 ttl=242 time=47.6 ms`<br>`^C`<br>After you have verified connectivity, press Ctrl+C to stop the pings. |

# Register the Firewall

| Register the Firewall | | |
|---|---|---|
| Step 1 | Log in to the web interface. | Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page. |
| Step 2 | Locate your serial number and copy it to the clipboard. | On the **Dashboard**, locate your **Serial Number** in the General Information section of the screen. |
| Step 3 | Go to the Palo Alto Networks Support site. | In a new browser tab or window, go to https://support.paloaltonetworks.com. |
| Step 4 | Register the device. The way you register depends on whether you already have a login to the support site. | • If this is the first Palo Alto Networks device you are registering and you do not yet have a login, click **Register** on the right side of the page. To register, you must provide your sales order number or customer ID, and the serial number of your firewall (which you can paste from your clipboard) or the authorization code you received with your order. You will also be prompted to set up a username and password for access to the Palo Alto Networks support community.<br> <br>• If you already have a support account, log in and then click **My Devices**. Scroll down to Register Device section at the bottom of the screen and enter the serial number of your firewall (which you can paste from your clipboard), your city and postal code and then click **Register Device**. |

 © Palo Alto Networks, Inc.

## Activate Licenses and Subscriptions

Before you can start using your firewall to secure the traffic on your network, you must activate the licenses for each of the services you purchased. Available licenses and subscriptions include the following:

- Threat Prevention—Provides antivirus, anti-spyware, and vulnerability protection.
- Decryption Mirroring—Provides the ability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as NetWitness or Solera—for archiving and analysis.
- URL Filtering—Allows you create security policy to enforce web access based on dynamic URL categories. You must purchase and install a subscription for one of the supported URL filtering databases: PAN-DB or BrightCloud. With PAN-DB, you can set up access to the PAN-DB public cloud or to the PAN-DB private cloud. For more information about URL filtering, see Control Access to Web Content.
- Virtual Systems—This license is required to enable support for multiple virtual systems on PA-2000 and PA-3000 Series firewalls. In addition, you must purchase a Virtual Systems license if you want to increase the number of virtual systems beyond the base number provided by default on PA-4000 Series, PA-5000 Series, and PA-7000 Series firewalls (the base number varies by platform). The PA-500, PA-200, and VM-Series firewalls do not support virtual systems.
- WildFire—Although basic WildFire support is included as part of the Threat Prevention license, the WildFire subscription service provides enhanced services for organizations that require immediate coverage for threats, frequent WildFire signature updates, advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet), as well as the ability to upload files using the WildFire API. A WildFire subscription is also required if your firewalls will be forwarding files to a WF-500 appliance.
- GlobalProtect—Provides mobility solutions and/or large-scale VPN capabilities. By default, you can deploy GlobalProtect portals and gateways (without HIP checks) without a license. If you want to use HIP checks, you will also need gateway licenses (subscription) for each gateway.

| **Activate Licenses** | | |
|---|---|---|
| Step 1 | Locate the activation codes for the licenses you purchased. | When you purchased your subscriptions you should have received an email from Palo Alto Networks customer service listing the activation code associated with each subscription. If you cannot locate this email, contact customer support to obtain your activation codes before you proceed. |
| Step 2 | Launch the web interface and go to the license page. | Select **Device > Licenses**. |

| Activate Licenses (Continued) | | |
|---|---|---|
| Step 3 | Activate each license you purchased. | After purchasing your licenses/subscriptions activate them in one of the following ways:<br>• **Retrieve license keys from license server**—Use this option if you activated your license on the support portal.<br>• **Activate feature using authorization code**—Use this option to enable purchased subscriptions using an authorization code for licenses that have not been previously activated on the support portal. When prompted, enter the **Authorization Code** and then click **OK**.<br>• **Manually upload license key**—Use this option if your device does not have connectivity to the Palo Alto Networks support site. In this case, you must download a license key file from the support site on an Internet connected computer and then upload to the device. |
| Step 4 | Verify that the license was successfully activated | On the **Device > Licenses** page, verify that the license was successfully activated. For example, after activating the WildFire license, you should see that the license is valid:<br> |
| Step 5 | (WildFire subscriptions only) Perform a commit to complete WildFire subscription activation. | After activating a WildFire subscription, a commit is required for the firewall to begin forwarding advanced file types:<br>• Commit any pending changes.<br>• Make a minor change and perform a commit. For example, update a rule description and commit the change.<br>Check that the WildFire Analysis profile rules include the advanced file types that are now supported with the WildFire subscription. If no change to any of the rules is required, make a minor edit to a rule description and perform a commit. |

 © Palo Alto Networks, Inc.

## Manage Content Updates

In order to stay ahead of the changing threat and application landscape, Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to the Palo Alto Networks devices. The devices access the web resources in the CDN to perform various App-ID and Content-ID functions. By default, the devices use the management port to access the CDN infrastructure for application updates, threat and antivirus signature updates, BrightCloud and PAN-DB database updates and lookups, and access to the Palo Alto Networks WildFire cloud. To ensure that you are always protected from the latest threats (including those that have not yet been discovered), you must ensure that you keep your devices up-to-date with the latest updates published by Palo Alto Networks.

The following content updates are available, depending on which subscriptions you have:

Although you can manually download and install content updates at any time, as a best practice you should Schedule each update. Scheduled updates occur automatically.

- **Antivirus**—Includes new and updated antivirus signatures, including signatures discovered by the WildFire cloud service. You must have a Threat Prevention subscription to get these updates. New antivirus signatures are published daily.
- **Applications**—Includes new and updated application signatures. This update does not require any additional subscriptions, but it does require a valid maintenance/support contract. New application updates are published weekly. To review the policy impact of new application updates, see Manage New App-IDs Introduced in Content Releases.
- **Applications and Threats**—Includes new and updated application and threat signatures. This update is available if you have a Threat Prevention subscription (and you get it instead of the Applications update). New Applications and Threats updates are published weekly.
- **GlobalProtect Data File**—Contains the vendor-specific information for defining and evaluating host information profile (HIP) data returned by GlobalProtect agents. You must have a GlobalProtect gateway license and create an update schedule in order to receive these updates.
- **BrightCloud URL Filtering**—Provides updates to the BrightCloud URL Filtering database only. You must have a BrightCloud subscription to get these updates. New BrightCloud URL database updates are published daily. If you have a PAN-DB license, scheduled updates are not required as devices remain in-sync with the servers automatically.
- **WildFire**—Provides near real-time malware and antivirus signatures created as a result of the analysis done by the WildFire cloud service. Without the subscription, you must wait 24 to 48 hours for the signatures to roll into the Applications and Threat update.

If your firewall does not have Internet access from the management port, you can download content updates from the Palo Alto Networks Support portal and then Upload them to your firewall.

If your firewall is deployed behind existing firewalls or proxy servers, access to these external resources might be restricted using access control lists that allow the firewall to only access a hostname or an IP address. In such cases, to allow access to the CDN, set the update server address to use the hostname `staticupdates.paloaltonetworks.com` or the IP address `199.167.52.15`. For details on setting up CDN access, see Content Delivery Network Infrastructure for Dynamic Updates.

| Update Content | |
|---|---|
| Step 1 Verify that the firewall points to the CDN infrastructure. | Select **Device > Setup > Services**.<br>• As a best practice, set the **Update Server** to access `updates.paloaltonetworks.com`. This allows the firewall to receive content updates from the server to which it is closest in the CDN infrastructure.<br>• (Optional) If the firewall has restricted access to the Internet, set the update server address to use the hostname `staticupdates.paloaltonetworks.com` or the IP address `199.167.52.15`.<br>• For additional security, select **Verify Update Server Identity**. The firewall verifies that the server from which the software or content package is download has an SSL certificate signed by a trusted authority. |
| Step 2 Launch the web interface and go to the Dynamic Updates page. | Select **Device > Dynamic Updates**. |
| Step 3 Check for the latest updates. | Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available:<br>• **Download**—Indicates that a new update file is available. Click the link to begin downloading the file directly to the firewall. After successful download, the link in the **Action** column changes from **Download** to **Install**.<br> <br>You cannot download the antivirus database until you have installed the Application and Threats database.<br>• **Upgrade**—Indicates that there is a new version of the BrightCloud database available. Click the link to begin the download and installation of the database. The database upgrade begins in the background; when completed a check mark displays in the **Currently Installed** column. Note that if you are using PAN-DB as your URL filtering database you will not see an upgrade link because the PAN-DB database automatically stays in sync with the server.<br> <br>To check the status of an action, click **Tasks** (on the lower right-hand corner of the window).<br>• **Revert**—Indicates that a previously installed version of the content or software version is available. You can choose to revert to the previously installed version. |

 © Palo Alto Networks, Inc.

| Update Content (Continued) | | |
|---|---|---|
| Step 4 | Install the updates.<br><br>Installation can take up to 20 minutes on a PA-200, PA-500, or PA-2000 Series device and up to two minutes on a PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7000 Series, or VM-Series firewall. | Click the **Install** link in the **Action** column. When the installation completes, a check mark displays in the **Currently Installed** column. |
| | | |
| Step 5 | Schedule each update.<br><br>Repeat this step for each update you want to schedule.<br><br>Stagger the update schedules because the firewall can only download one update at a time. If you schedule the updates to download during the same time interval, only the first download will succeed. | 1. Set the schedule of each update type by clicking the **None** link.<br> <br>2. Specify how often you want the updates to occur by selecting a value from the **Recurrence** drop-down. The available values vary by content type (WildFire updates are available **Every 15 minutes**, **Every 30 minutes** or **Every Hour** whereas all other content types can be scheduled for **Daily** or **Weekly** update).<br>3. Specify the **Time** and (or, minutes past the hour in the case of WildFire), if applicable depending on the **Recurrence** value you selected, **Day** of the week that you want the updates to occur.<br>4. Specify whether you want the system to **Download Only** or, as a best practice, **Download And Install** the update.<br>5. In rare instances, errors in content updates may be found. For this reason, you may want to delay installing new updates until they have been released for a certain number of hours. You can specify how long after a release to wait before performing a content update by entering the number of hours to wait in the **Threshold (Hours)** field.<br>6. Click **OK** to save the schedule settings.<br>7. Click **Commit** to save the settings to the running configuration. |

## Install Software Updates

When installing a new firewall, it is a good idea to upgrade to the latest software update (or to the update version that your reseller or Palo Alto Networks Systems Engineer recommends) to take advantage of the latest fixes and security enhancements. Before updating the software, make sure you have the latest content updates as detailed in Manage Content Updates (the Release Notes for a software update specify the minimum content release version supported in the release).

 © Palo Alto Networks, Inc.

# Create the Security Perimeter

Traffic must pass through the firewall in order for the firewall to manage and control it. Physically, traffic enters and exits the firewall through *interfaces*. The firewall determines how to act on a packet based on whether the packet matches a *security policy rule*. At the most basic level, each security policy rule must identify where the traffic came from and where it is going. On a Palo Alto Networks next-generation firewall, security policy rules are applied between zones. A *zone* is a grouping of interfaces (physical or virtual) that provides an abstraction for an area of trust for simplified policy enforcement. For example, in the following topology diagram, there are three zones: Trust, Untrust, and DMZ. Traffic can flow freely within a zone, but traffic will not be able to flow between zones until you define a security policy rule that allows it.

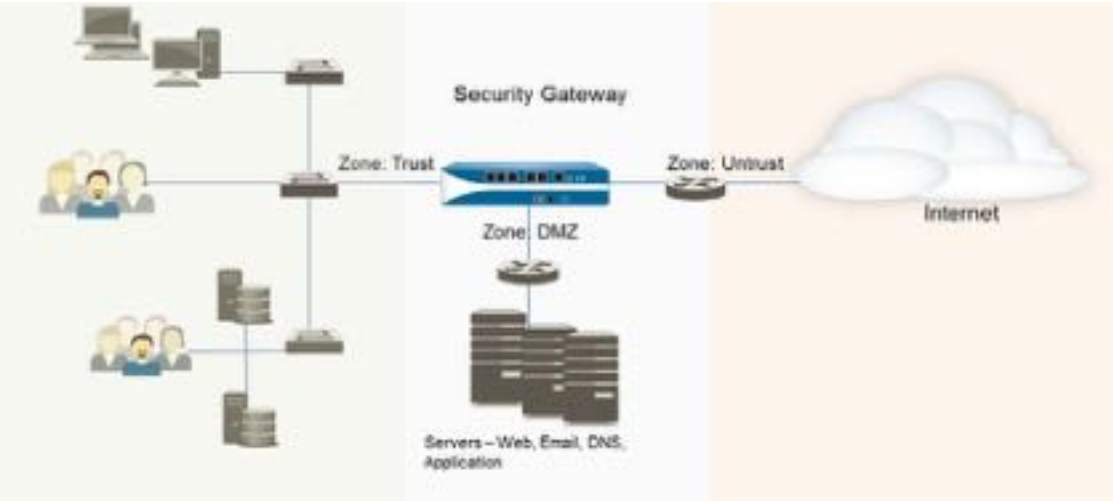


The following topics describe the components of the security perimeter and provide steps for configuring the firewall interfaces, defining zones, and setting up a basic security policy that allows traffic from your internal zone to the Internet and to the DMZ. By initially creating a basic security policy rulebase like this, you will be able to analyze the traffic running through your network and use this information to define more granular policies for safely enabling applications while preventing threats.

- Basic Interface Deployments
- About Security Policy
- Plan the Deployment
- Configure Interfaces and Zones
- Set Up Basic Security Policies

If you use private IP addresses in your internal networks, you will also need to configure network address translation (NAT).

# Basic Interface Deployments

All Palo Alto Networks next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, enabling you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports. The following sections provide basic information on each type of deployment.

- Virtual Wire Deployments
- Layer 2 Deployments
- Layer 3 Deployments

For more detailed deployment information, refer to Designing Networks with Palo Alto Networks Firewalls.

## Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together. By using a virtual wire, you can install the firewall in any network environment without reconfiguring adjacent devices. If necessary, a virtual wire can block or allow traffic based on the virtual LAN (VLAN) tag values. You can also create multiple subinterfaces and classify traffic according to an IP Address (address, range, or subnet), VLAN, or a combination of the two.

By default, the virtual wire (named **default-vwire**) binds Ethernet ports 1 and 2 and allows all untagged traffic. Choose this deployment to simplify installation and configuration and/or avoid configuration changes to surrounding network devices.

A virtual wire is the default configuration, and should be used only when no switching or routing is needed. If you do not plan to use the default virtual wire, you should manually delete the configuration before proceeding with interface configuration to prevent it from interfering with other interface settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up a Data Port for Access to External Services.

## Layer 2 Deployments

In a Layer 2 deployment, the firewall provides switching between two or more interfaces. Each group of interfaces must be assigned to a VLAN object in order for the firewall to switch between them. The firewall will perform VLAN tag switching when Layer 2 subinterfaces are attached to a common VLAN object. Choose this option when switching is required.

For more information on Layer 2 deployments, refer to the Layer 2 Networking Tech Note and/or the Securing Inter VLAN Traffic Tech Note.

 © Palo Alto Networks, Inc.

## Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between ports. An IP address must be assigned to each interface and a virtual router must be defined to route the traffic. Choose this option when routing is required.

You must assign an IP address to each physical Layer 3 interface you configure. You can also create logical subinterfaces for each physical Layer 3 interface that allows you to segregate the traffic on the interface based on VLAN tag (when VLAN trunking is in use) or by IP address, for example for multi-tenancy.

In addition, because the firewall must route traffic in a Layer 3 deployment, you must configure a virtual router. You can configure the virtual router to participate with dynamic routing protocols (BGP, OSPF, or RIP) as well as adding static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that are not shared between virtual routers, enabling you to configure different routing behaviors for different interfaces. For more information on routing integrations on the firewall, see the PAN-OS Admin Guide.

The configuration example in this chapter illustrates how to integrate the firewall into your Layer 3 network using static routes.

# About Security Policy

Security Policy protects network assets from threats and disruptions and aids in optimally allocating network resources for enhancing productivity and efficiency in business processes. On the Palo Alto Networks firewall, security policy rules determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service.

For traffic that doesn't match any defined rules, the default rules apply. The default rules—displayed at the bottom of the security rulebase—are predefined to allow all intrazone (within the zone) traffic and deny all interzone (between zones) traffic. Although these rules are part of the pre-defined configuration and are read-only by default, you can override them and change a limited number of settings, including the tags, action (allow or deny), log settings, and security profiles.

Security policies rules are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log, if logging is enabled for that rule. The logging options are configurable for each rule, and can for example be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.

- Actions in Security Policy
- About Policy Objects
- About Security Profiles

## Actions in Security Policy

For traffic that matches the attributes defined in a security policy, you can apply the following actions:

| Action | Description |
|---|---|
| **Allow** (default action) | Allows the traffic. |
| **Deny** | Blocks traffic, and enforces the default *Deny Action* defined for the application that is being denied. To view the deny action defined by default for an application, view the application details in **Objects > Applications** or check the application details in Applipedia. |
| **Drop** | Silently drops the traffic; for an application, it overrides the default deny action. A TCP reset is not sent to the host/application.<br>For Layer 3 interfaces, to optionally send an ICMP unreachable response to the client, set Action: **Drop** and enable the **Send ICMP Unreachable** check box. When enabled, the firewall sends the ICMP code for *communication with the destination is administratively prohibited*— ICMPv4: Type 3, Code 13; ICMPv6: Type 1, Code 1. |
| **Reset client** | Sends a TCP reset to the client-side device. |
| **Reset server** | Sends a TCP reset to the server-side device. |
| **Reset both** | Sends a TCP reset to both the client-side and server-side devices. |

 © Palo Alto Networks, Inc.

A reset is sent only after a session is formed. If the session is blocked before a 3-way handshake is completed, the firewall will not send the reset.

For a TCP session with a reset action, the firewall does not send an ICMP Unreachable response.

For a UDP session with a drop or reset action, if the **ICMP Unreachable** check box is selected, the firewall sends an ICMP message to the client.

## About Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With Policy Objects that are a collective unit, you can reference the object in security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

Some examples of address and application policy objects are shown in the security policies that are included in Create Security Rules. For information on the other policy objects, see Enable Basic Threat Prevention Features.

## About Security Profiles

While security policies enable you to allow or deny traffic on your network, security profiles help you define an *allow but scan* rule, which scan allowed applications for threats. When traffic matches the allow rule defined in the security policy, the Security Profiles that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.

Security profiles are not used in the match criteria of a traffic flow. The security profile is applied to scan traffic after the application or category is allowed by the security policy.

The different types of security profiles that can be attached to security policies are: Antivirus, Anti-Spyware, Vulnerability Protection, URL Filtering, File Blocking, and Data Filtering. The firewall provides default security profiles that you can use out of the box to begin protecting your network from threats. See Create Security Rules for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, you can create custom profiles. See Scan Traffic for Threats for more information.

 © Palo Alto Networks, Inc.

## Plan the Deployment

Before you begin configuring interfaces and zones, take some time to plan the zones you will need based on the different usage requirements within your organization. In addition, you should gather all of the configuration information you will need ahead of time. At a basic level, you should plan which interfaces will belong to which zones. For Layer 3 deployments you'll also need to obtain the required IP addresses and network configuration information from your network administrator, including information on how to configure the routing protocol or static routes required for the virtual router configuration. The example in this chapter will be based on the following topology:

**Figure: Layer 3 Topology Example**

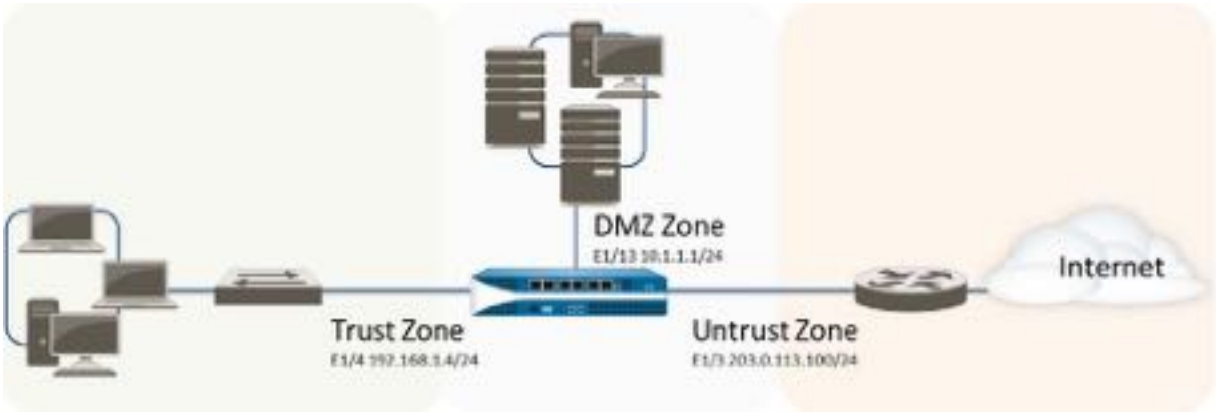


The following table shows the information we will use to configure the Layer 3 interfaces and their corresponding zones as shown in the sample topology.

| Zone | Deployment Type | Interface(s) | Configuration Settings |
|---|---|---|---|
| Untrust | L3 | Ethernet1/3 | **IP address:** 203.0.113.100/24<br>**Virtual router:** default<br>**Default route:** 0.0.0.0/0<br>**Next hop:** 203.0.113.1 |
| Trust | L3 | Ethernet1/4 | **IP address:** 192.168.1.4/24<br>**Virtual router:** default |
| DMZ | L3 | Ethernet1/13 | **IP address:** 10.1.1.1/24<br>**Virtual router:** default |

## Configure Interfaces and Zones

After you plan your zones and the corresponding interfaces, you can configure them on the device. The way you configure each interface depends on your network topology.

The following procedure shows how to configure a Layer 3 deployment as depicted in Figure: Layer 3 Topology Example.

The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and virtual router). If you do not plan to use the default virtual wire, you must manually delete the configuration and commit the change before proceeding to prevent it from interfering with other settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up a Data Port for Access to External Services.

| **Set Up Interfaces and Zones** | | |
|---|---|---|
| Step 1 | Configure a default route to your Internet router. | 1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog.<br>2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and enter the route in the **Destination** field (for example, 0.0.0.0/0).<br>3. Select the **IP Address** radio button in the **Next Hop** field and then enter the IP address and netmask for your Internet gateway (for example, 203.00.113.1).<br>4. Click **OK** twice to save the virtual router configuration. |
| Step 2 | Configure the external interface (the interface that connects to the Internet). | 1. Select **Network > Interfaces** and then select the interface you want to configure. In this example, we are configuring Ethernet1/3 as the external interface.<br>2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.<br>3. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Untrust, and then click **OK**.<br>4. In the **Virtual Router** drop-down, select **default**.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 208.80.56.100/24.<br>6. To enable you to ping the interface, select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. Enter a **Name** for the profile, select **Ping** and then click **OK**.<br>7. To save the interface configuration, click **OK**. |

 © Palo Alto Networks, Inc.

| Set Up Interfaces and Zones (Continued) | | |
|---|---|---|
| Step 3 | Configure the interface that connects to your internal network.<br><br>In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT. | 1. Select **Network > Interfaces** and select the interface you want to configure. In this example, we are configuring Ethernet1/4 as the internal interface.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Trust, and then click **OK**.<br>4. Select the same Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 4 | Configure the interface that connects to the DMZ. | 1. Select the interface you want to configure.<br>2. Select **Layer3** from the **Interface Type** drop-down. In this example, we are configuring Ethernet1/13 as the DMZ interface.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example DMZ, and then click **OK**.<br>4. Select the Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 5 | Save the interface configuration. | Click **Commit**. |
| Step 6 | Cable the firewall. | Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment. |
| Step 7 | Verify that the interfaces are active. | From the web interface, select **Network > Interfaces** and verify that icon in the Link State column is green. You can also monitor link state from the **Interfaces** widget on the **Dashboard**.<br><br> 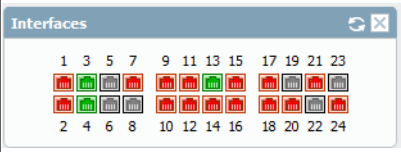 |

# Set Up Basic Security Policies

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Captive Portal, Denial of Service, and Zone protection policies. All these different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. This section covers basic security policies and the default security profiles:

- Create Security Rules
- Test Your Security Policies
- Monitor Network Traffic

## Create Security Rules

Security policies reference security zones and enable you to allow, restrict, and track traffic on your network. Because each zone implies a level of trust, the implicit rule for passing traffic between two different zones is deny, and the traffic within a zone is permitted. To allow traffic between two different zones, you must create a security rule that allows traffic to flow between them.

While setting up the basic framework for securing the enterprise perimeter, it's good idea to start with a simple security policy that allows traffic between the different zones without being too restrictive. As illustrated in the following section, our objective is to minimize the likelihood of breaking applications that users on the network need access to, while providing visibility into the applications and the potential threats for your network.

When defining policies make sure that you do not create a policy that denies all traffic from *any* source zone to *any* destination zone as this will break intra-zone traffic that is implicitly allowed. By default, intra-zone traffic is permitted because the source and destination zones are the same and therefore share the same level of trust.

 © Palo Alto Networks, Inc.

## Define Basic Security Rules

**Step 1** Permit Internet access for all users on the enterprise network.

Zone: Trust to Untrust

By default, the firewall includes a security rule named *rule1* that allows all traffic from Trust zone to Untrust zone. You can either delete the rule or modify the rule to reflect your zone-naming convention.

To safely enable applications that are required for day-to-day business operations we will create a simple rule that allows access to the Internet. To provide basic threat protection, we will attach the default security profiles available on the firewall.

1. Select **Policies > Security** and click **Add.**
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, Set the **Destination Zone** to Untrust.

   To scan policy rules and visually identify the zones on each rule, create a tag with the same name as the zone. For example, to color code the Trust zone as green, select **Objects > Tags,** click **Add** and **Name** the tag Trust, and select the **Color** green.

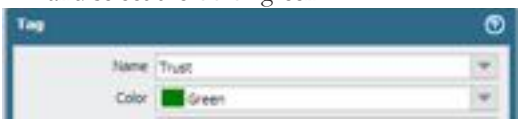


5. In the **Service/ URL Category** tab, select service-http and service-https.
6. In the **Actions** tab, complete these tasks:
   a. Set the **Action Setting** to **Allow**.
   b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection and URL filtering, under **Profile Setting.**
7. Verify that logging is enabled at the end of a session under **Options**. Only traffic that matches a security rule will be logged.

**Step 2** Permit users on the internal network to access the servers in the DMZ.

Zone: Trust to DMZ

If using IP addresses for configuring access to the servers in the DMZ, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT addresses), and the post-NAT zone.

1. Click **Add** in the **Policies > Security** section.
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, set the **Destination Zone** to DMZ.
5. In the **Service/ URL Category** tab, make sure the **Service** is set to **application-default**.
6. In the **Actions** tab, set the **Action Setting** to Allow.
7. Leave all the other options at the default values.

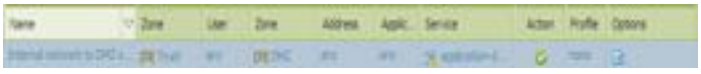

**Define Basic Security Rules (Continued)**

**Step 3** Restrict access from the Internet to the servers on the DMZ to specific server IP addresses only.

For example, you might only allow users to access the webmail servers from outside.

Zone: Untrust to DMZ

To restrict inbound access to the DMZ from the Internet, configure a rule that allows access only to specific servers IP addresses and on the default ports that the applications use.

1. Click **Add** to add a new rule, and give it a descriptive name.
2. In the **Source** tab, set the **Source Zone** to Untrust.
3. In the **Destination** tab, set the **Destination Zone** to DMZ.
4. Set the **Destination Address** to the **Public web server** address object you created earlier. The public web server address object references the public IP address—208.80.56.11/24—of the web server that is accessible on the DMZ.
5. Select the webmail application in the **Application** tab.

   The **Service** is set to **application-default** by default.

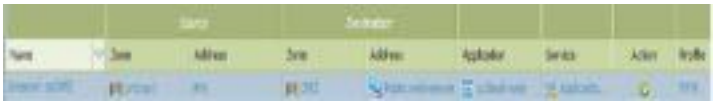

6. Set the **Action Setting** to **Allow**.

**Step 4** Allow access from the DMZ to your internal network (Trust zone). To minimize risk, you will allow traffic only between specific servers and destination addresses. For example, if you have an application server on the DMZ that needs to communicate with a specific database server in your Trust zone, create a rule to allow traffic between a specific source to a specific destination.

Zone: DMZ to Trust

1. Click **Add** to add a new rule, and give it a descriptive name.
2. Set the **Source Zone** to DMZ.
3. Set the **Destination Zone** to Trust.
4. Create a an address object that specifies the server(s) on your Trust zone that can be accessed from the DMZ.

5. In the **Destination** tab on the Security Policy rule, set the **Destination Address** to the Address object you created above.
6. In the **Actions** tab, complete these tasks:
   a. Set the **Action Setting** to **Allow**.
   b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection, under **Profile Setting.**
   c. In the Other Settings section, select the option to **Disable Server Response Inspection**. This setting disables the antivirus and anti-spyware scanning on the server-side responses, and thus reduces the load on the firewall.

 © Palo Alto Networks, Inc.

| Define Basic Security Rules (Continued) | | |
|---|---|---|
| Step 5 | Enable the servers on the DMZ to obtain updates and hot fixes from the Internet. Say, for example, you would like to allow the Microsoft Update service.<br><br>Zone: DMZ to Untrust | 1. Add a new rule and give it a descriptive label.<br>2. Set the **Source Zone** to DMZ.<br>3. Set the **Destination Zone** to Untrust.<br>4. Create an application group to specify the applications that you would like to allow. In this example, we allow Microsoft updates (ms-updates) and dns.<br> 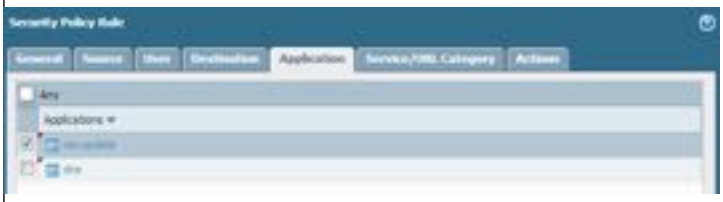 <br>The **Service** is set to **application-default** by default. This allows the firewall to permit the applications only when they use the standard ports associated with these applications.<br>5. Set the **Action Setting** to **Allow**.<br>6. Attach the default profiles for antivirus, anti-spyware, and vulnerability protection, under **Profiles**.<br> |
| Step 6 | Save your policies to the running configuration on the device. | Click **Commit**. |

## Test Your Security Policies

To verify that you have set up your basic policies effectively, test whether your security policies are being evaluated and determine which security rule applies to a traffic flow.

| Verify Policy Match Against a Flow | |
|---|---|
| To verify the policy rule that matches a flow, use the following CLI command:<br>**test security-policy-match source** <IP_address> **destination** <IP_address> **destination port** <port_number> **protocol** <protocol_number><br>The output displays the best rule that matches the source and destination IP address specified in the CLI command. | For example, to verify the policy rule that will be applied for a server on the DMZ with the IP address 208.90.56.11 when it accesses the Microsoft update server, you will try the following command:<br>`test security-policy-match source 208.80.56.11 destination 176.9.45.70 destination-port 80 protocol 6`<br>`"Updates-DMZ to Internet" { from dmz; source any; source-region any; to untrust; destination any; destination-region any; user any; category any; application/service[dns/tcp/any/53 dns/udp/any/53 dns/udp/any/5353 ms-update/tcp/any/80 ms-update/tcp/any/443]; action allow; terminal yes;` |

## Monitor Network Traffic

Now that you have a basic security policy, you can review the statistics and data in the Application Command Center (ACC), traffic logs, and the threat logs to observe trends on your network, to identify where you need to create more granular policies.

| Monitor Network Traffic | |
|---|---|
| • Use the Application Command Center and Use the Automated Correlation Engine. | In the ACC, review the most used applications and the high-risk applications on your network. The ACC graphically summarizes the log information to highlight the applications traversing the network, who is using them (with User-ID enabled), and the potential security impact of the content to help you identify what is happening on the network in real time. You can then use this information to create appropriate security policy rules that block unwanted applications, while allowing and enabling applications in a secure manner.<br>The Compromised Hosts widget in **ACC > Threat Activity** displays potentially compromised hosts on your network and the logs and match evidence that corroborates the events. |

 © Palo Alto Networks, Inc.

| Monitor Network Traffic | |
|---|---|
| • Determine what updates/modifications are required for your network security policy rules and implement the changes. | For example:<br>• Evaluate whether to allow content based on schedule, users, or groups.<br>• Allow or control certain applications or functions within an application.<br>• Decrypt and inspect content.<br>• Allow but scan for threats and exploits.<br>For information on refining your security policies and for attaching custom security profiles, see Enable Basic Threat Prevention Features. |
| • View the Log Files. | Specifically, view the traffic and threat logs (**Monitor > Logs**).<br> Traffic logs are dependent on how your security policies are defined and set up to log traffic. The Application Usage widget in the **ACC**, however, records applications and statistics regardless of policy configuration; it shows all traffic that is allowed on your network, therefore it includes the inter-zone traffic that is allowed by policy and the same zone traffic that is allowed implicitly |
| • Monitor Web Activity of Network Users. | Review the URL filtering logs to scan through alerts, denied categories/URLs. URL logs are generated when a traffic matches a security rule that has a URL filtering profile attached with an action of alert, continue, override or block. |

# Enable Basic Threat Prevention Features

The Palo Alto Networks next-generation firewall has unique threat prevention capabilities that allow it to protect your network from attack despite evasive, tunneled, or circumvention techniques. The threat prevention features on the firewall include the WildFire service, the Security Profiles that support Antivirus, Anti-Spyware, Vulnerability Protection, URL Filtering, File Blocking and Data Filtering capabilities and the Denial of Service (DoS) and Zone protection functionality.

Before you can apply threat prevention features, you must first configure zones—to identify one or more source or destination interfaces—and security policy rules. To configure interfaces, zones, and the policies that are needed to apply threat prevention features, see Configure Interfaces and Zones and Set Up Basic Security Policies.

To begin protecting your network from threats start here:

- Enable WildFire
- Scan Traffic for Threats
- Control Access to Web Content

 © Palo Alto Networks, Inc.

## Enable WildFire

The WildFire service is included as part of the base product. The WildFire service enables the firewall to forward attachments to a sandbox environment where applications are run to detect any malicious activity. As new malware is detected by the WildFire service, malware signatures are automatically generated and are made available within 24-48 hours in the antivirus daily downloads. Your threat prevention subscription entitles you to antivirus signature updates that include signatures discovered by WildFire.

Consider purchasing the WildFire subscription service for these additional benefits:

- Sub-hourly (as often as every 15 minutes) WildFire signature updates.
- Advanced file type forwarding for APKs, Flash files, PDFs, Microsoft Office files, Java Applets, Java files (.jar and .class), and HTTP/HTTPS email links contained in SMTP and POP3 email messages. Portable Executable (PE) files can be forwarded for WildFire analysis both with and without a WildFire subscription.
- Ability to upload files using the WildFire API.
- Ability to forward files to a private WF-500 WildFire appliance. When using a WildFire appliance, you can also set up a WildFire hybrid cloud, enabling the WildFire appliance to analyze sensitive file types locally, while other files less sensitive file types (such as PEs) or file types not supported for WildFire appliance analysis (such as APKs) can be forwarded to the WildFire cloud.

| **Enable WildFire** | | |
|---|---|---|
| Step 1 | Confirm that your device is registered and that you have a valid support account as well as any subscriptions you require. | 1. Go to the Palo Alto Networks Support Site, log in, and select **My Devices**.<br>2. Verify that the firewall is listed. If it is not listed, see Register the Firewall.<br>3. (Optional) Activate Licenses and Subscriptions. |
| Step 2 | Set the WildFire forwarding options.<br> If you do not have a WildFire subscription you can only forward PEs. | 1. Select **Device > Setup > WildFire** and edit the General Settings.<br>2. (Optional) Specify the WildFire cloud or WildFire appliance (or both) to which the firewall will forward files for analysis. By default, the firewall will forward files to the public WildFire cloud hosted in the United States (wildfire.paloaltonetworks.com). To forward files to the WildFire cloud hosted in Japan or to enable file-forwarding to a WildFire appliance, update the following fields:<br>• **WildFire Public Cloud**—To forward files to the public WildFire cloud running in Japan, enter `wildfire.paloaltonetworks.jp`.<br>• **WildFire Private Cloud**—To forward files to a private WildFire cloud, enter the IP address or FQDN of your WF-500 WildFire appliance.<br>3. (Optional) If you want to change the maximum file size that the firewall can forward for a specific type of file, modify the value in the corresponding field.<br>4. Click **OK** to save your changes. |

| Enable WildFire (Continued) | | |
|---|---|---|
| Step 3 | Set up a WildFire Analysis profile to forward files to WildFire. | 1. Select **Objects > Security Profiles > WildFire Analysis** and click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the profile.<br>3. Click **Add** to create a forwarding rule and enter a name.<br>4. Define traffic to be forwarded to the WildFire service based on **Applications**, **File Types,** or transmission **Direction**.<br>5. In the **Analysis** column, select **public-cloud** to forward the defined files to the WildFire cloud or select **private-cloud** to forward files to the WildFire appliance.<br>6. (Optional) Add additional forwarding rules as necessary.<br>7. Click **OK** to save the profile. |
| Step 4 | Attach the WildFire Analysis profile to the security policies that allow access to the Internet. | 1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Create Security Rules.<br>2. Click the **Actions** tab within the security policy.<br>3. In the Profile Settings section, click the drop-down and select the WildFire Analysis profile you created for WildFire forwarding. (If you don't see a drop-down for selecting a profile, select **Profiles** from the **Profile Type** drop-down. |
| Step 5 | Save the configuration. | Click **Commit**. |
| Step 6 | Verify that the firewall is forwarding files to the WildFire cloud or the WildFire appliance. | 1. Select **Monitor > Logs > WildFire Submissions** to WildFire logs. For each log entry displayed, the firewall has successfully forwarded the file to WildFire and WildFire has returned a file analysis report.<br>2. Check the **WildFire Cloud** column to view if a file was forwarded to the WildFire cloud or the WildFire appliance for analysis. |

 © Palo Alto Networks, Inc.

# Scan Traffic for Threats

Security Profiles provide threat protection in security policies. For example, you can apply an antivirus profile to a security policy and all traffic that matches the security policy will be scanned for viruses.

The following sections provide steps for setting up a basic threat prevention configuration:

- Set Up Antivirus, Anti-Spyware, and Vulnerability Protection
- Set Up File Blocking

## Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

Every Palo Alto Networks next-generation firewall comes with predefined Antivirus, Anti-Spyware, and Vulnerability Protection profiles that you can attach to security policies. There is one predefined Antivirus profile, **default**, which uses the default action for each protocol (block HTTP, FTP, and SMB traffic and alert on SMTP, IMAP, and POP3 traffic). There are two predefined Anti-Spyware and Zone Protection profiles:

- **default**—Applies the default action to all client and server critical, high, and medium severity spyware/vulnerability protection events. It does not detect low and informational events.
- **strict**—Applies the block response to all client and server critical, high and medium severity spyware/vulnerability protection events and uses the default action for low and informational events.

To ensure that the traffic entering your network is free from threats, attach the predefined profiles to your basic web access policies. As you monitor the traffic on your network and expand your policy rulebase, you can then design more granular profiles to address your specific security needs.

| **Set up Antivirus/Anti-Spyware/Vulnerability Protection** | | |
|---|---|---|
| Step 1 | Verify that you have a Threat Prevention license. | • The Threat Prevention license bundles the Antivirus, Anti-Spyware, and the Vulnerability Protection features in one license.<br>• Select **Device > Licenses** to verify that the **Threat Prevention** license is installed and valid (check the expiration date). |
| Step 2 | Download the latest antivirus threat signatures. | 1. Select **Device > Dynamic Updates** and click **Check Now** at the bottom of the page to retrieve the latest signatures.<br>2. In the **Actions** column, click **Download** to install the latest Antivirus, and Applications and Threats signatures. |

**Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued)**

Step 3 Schedule signature updates.

Perform a **download-and-install** on a daily basis for antivirus updates and weekly for applications and threats updates.

1. From **Device > Dynamic Updates**, click the text to the right of **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.
2. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select **Download Only**, you would need to manually go in and click the **Install** link in the **Action** column to install the signature. When you click **OK**, the update is scheduled. No commit is required.
3. (Optional) You can also enter the number of hours in the **Threshold** field to indicate the minimum age of a signature before a download will occur. For example, if you entered **10**, the signature must be at least 10 hours old before it will be downloaded, regardless of the schedule.
4. In an HA configuration, you can also click the **Sync To Peer** option to synchronize the content update with the HA peer after download/install. This will not push the schedule settings to the peer device, you need to configure the schedule on each device.

Recommendations for HA Configurations:

- **Active/Passive HA**—If the MGT port is used for antivirus signature downloads, you should configure a schedule on both devices and both devices will download/install independently. If you are using a data port for downloads, the passive device will not perform downloads while it is in the passive state. In this case you would set a schedule on both devices and then select the **Sync To Peer** option. This will ensure that whichever device is active, the updates will occur and will then push to the passive device.
- **Active/Active HA**—If the MGT port is used for antivirus signature downloads on both devices, then schedule the download/install on both devices, but do not select the **Sync To Peer** option. If you are using a data port, schedule the signature downloads on both devices and select **Sync To Peer**. This will ensure that if one device in the active/active configuration goes into the active-secondary state, the active device will download/install the signature and will then push it to the active-secondary device.

 © Palo Alto Networks, Inc.

| Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued) | |
|---|---|
| Step 4 Attach the security profiles to a security policy.<br><br>Attach a clone of a predefined security profile to your basic security policies. That way, if you want to customize the profile you can do so without deleting the read-only predefined **strict** or **default** profile and attaching a customized profile. | 1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab.<br>2. In **Profile Settings**, click the drop-down next to each security profile you would like to enable. In this example we choose default for **Antivirus, Vulnerability Protection,** and **Anti-Spyware**.<br><br>If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.<br><br> 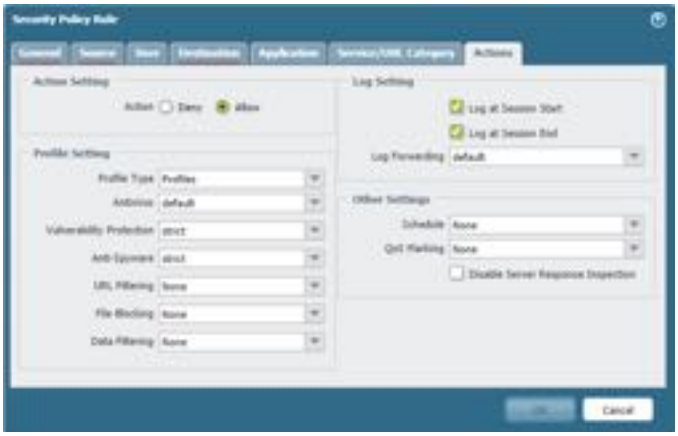 |
| Step 5 Save the configuration. | Click **Commit**. |

## Set Up File Blocking

File Blocking Profiles allow you to identify specific file types that you want to want to block or monitor. The following workflow shows how to set up a basic file blocking profile that prevents users from downloading executable files from the Internet.

| Configure File Blocking | |
|---|---|
| Step 1 Create the file blocking profile. | 1. Select **Objects > Security Profiles > File Blocking** and click **Add**.<br>2. Enter a **Name** for the file blocking profile, for example `Block_EXE`.<br>3. Optionally enter a **Description**, such as `Block users from downloading exe files from websites`. |
| Step 2 Configure the file blocking options. | 1. Click **Add** to define the profile settings.<br>2. Enter a **Name**, such as `BlockEXE`.<br>3. Set the **Applications** to which to apply file blocking, or leave it set to **any**.<br>4. Set **File Types** to block. For example, to block download of executables, you would select **exe**.<br>5. Specify the **Direction** in which to block files: **download**, **upload**, or **both**.<br>6. Set the **Action** to one of the following:<br>• **continue**—(web traffic only) Files matching the selected criteria will trigger a customizable response page that requires users to click **Continue** in order to proceed with the download/upload. You must enable response pages on the associated interfaces if you plan to use this option (Step 4).<br>• **block**—Files matching the selected criteria will be blocked from download/upload.<br>• **alert**—Files matching the selected criteria will be allowed, but will generate a log entry in the data filtering log.<br> <br>7. Click **OK** to save the profile. |
| Step 3 Attach the file blocking profile to the security policies that allow access to content. | 1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Create Security Rules.<br>2. Click the **Actions** tab within the security policy.<br>3. In the Profile Settings section, click the drop-down and select the file blocking profile you created.<br>If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down. |

 © Palo Alto Networks, Inc.

| Configure File Blocking (Continued) | | |
|---|---|---|
| Step 4 | Enable response pages in the management profile for each interface on which you are attaching file blocking profile with a **continue** action. | 1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.<br>2. Select **Response Pages**, as well as any other management services required on the interface.<br>3. Click **OK** to save the interface management profile.<br>4. Select **Network > Interfaces** and select the interface to which to attach the profile.<br>5. On the **Advanced > Other Info** tab, select the interface management profile you just created.<br>6. Click **OK** to save the interface settings. |
| Step 5 | Test the file blocking configuration. | Access a client PC in the trust zone of the firewall and attempt to download an.exe file from a website in the untrust zone. Make sure the file is blocked as expected based on the action you defined in the file blocking profile:<br>• If you selected **alert** as the action, check the data filtering log to make sure you see a log entry for the request.<br>• If you selected **block** as the action, the File Blocking Block Page response page should display.<br>• If you selected the **continue** action, the File Blocking Continue Page response page should display. Click **Continue** to download the file. The following shows the default File Blocking Continue Page.<br> |

## Control Access to Web Content

URL Filtering provides visibility and control over web traffic on your network. With URL filtering enabled, the firewall can categorize web traffic into one or more (from approximately 60) categories. You can then create policies that specify whether to allow, block, or log (alert) traffic based on the category to which it belongs. The following workflow shows how to enable PAN-DB for URL filtering, create security profiles, and attach them to security policies to enforce a basic URL filtering policy.

| **Configure URL Filtering** | | |
|---|---|---|
| Step 1 | Confirm license information for URL Filtering. | 1. Obtain and install a URL Filtering license. See Activate Licenses and Subscriptions for details.<br>2. Select **Device > Licenses** and verify that the URL Filtering license is valid.<br> |
| Step 2 | Download the seed database and activate the license. | 1. To download the seed database, click **Download** next to **Download Status** in the PAN-DB URL Filtering section of the Licenses page.<br>2. Choose a region (North America, Europe, APAC, Japan) and then click **OK** to start the download.<br>3. After the download completes, click **Activate**.<br> |
| Step 3 | Create a URL filtering profile.<br>Because the default URL filtering profile blocks risky and threat-prone content, clone this profile when creating a new profile in order to preserve the default settings. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select the default profile and then click **Clone**. The new profile will be named default-1.<br>3. Select the new profile and rename it. |

 © Palo Alto Networks, Inc.

| Configure URL Filtering (Continued) | |
|---|---|
| Step 4 Define how to control access to web content.<br><br>If you are not sure what traffic you want to control, consider setting the categories (except for those blocked by default) to alert. You can then use the visibility tools on the firewall, such as the ACC and App Scope, to determine which web categories to restrict to specific groups or to block entirely. You can then go back and modify the profile to block and allow categories as desired.<br><br>You can also define specific sites to always allow or always block regardless of category and enable the safe search option to filter search results when defining the URL Filtering profile. | 1. For each category that you want visibility into or control over, select a value from the **Action** column as follows:<br>• If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **allow**.<br>• For visibility into traffic to sites in a category, select **alert**.<br>• To present a response page to users attempting to access a particular category to alert them to the fact that the content they are accessing might not be work appropriate, select **continue**.<br>• To prevent access to traffic that matches the associated policy, select **block** (this also generates a log entry).<br> <br>2. Click **OK** to save the URL filtering profile. |
| Step 5 Attach the URL filtering profile to a security policy. | 1. Select **Policies > Security**.<br>2. Select the desired policy to modify it and then click the **Actions** tab.<br>3. If this is the first time you are defining a security profile, select **Profiles** from the **Profile Type** drop-down.<br>4. In the **Profile Settings** list, select the profile you just created from the **URL Filtering** drop-down. (If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.)<br>5. Click **OK** to save the profile.<br>6. **Commit** the configuration. |

| Configure URL Filtering (Continued) | | |
|---|---|---|
| Step 6 | Enable Response Pages in the management profile for each interface on which you are filtering web traffic. | 1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.<br>2. Select **Response Pages**, as well as any other management services required on the interface.<br>3. Click **OK** to save the interface management profile.<br>4. Select **Network > Interfaces** and select the interface to which to attach the profile.<br>5. On the **Advanced > Other Info** tab, select the interface management profile you just created.<br>6. Click **OK** to save the interface settings. |
| Step 7 | Save the configuration. | Click **Commit**. |
| Step 8 | Test the URL filtering configuration. | Access a client PC in the trust zone of the firewall and attempt to access a site in a blocked category. Make sure URL filtering is applied based on the action you defined in the URL filtering profile:<br>• If you selected **alert** as the action, check the data filtering log to make sure you see a log entry for the request.<br>• If you selected the **continue** action, the URL Filtering Continue and Override Page response page should display. **Continue** to the site.<br>• If you selected **block** as the action, the URL Filtering and Category Match Block Page response page should display as follows:<br> |

## For More Information

For more detailed information on how to protect your enterprise from threats, see Threat Prevention. For details on how to scan encrypted (SSH or SSL) traffic for threats, see Decryption.

For information about the threats and applications that Palo Alto Networks products can identify, visit the following links:

- Applipedia—Provides details on the applications that Palo Alto Networks can identify.
- Threat Vault—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

 © Palo Alto Networks, Inc.

# Best Practices for Completing the Firewall Deployment

Now that you have integrated the firewall into your network and enabled the basic security features, you can begin configuring more advanced features. Here are some things to consider next:

- Learn about the different Management Interfaces that are available to you and how to access and use them.
- Set up High Availability—High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in a two-device cluster provides redundancy and allows you to ensure business continuity.
- Configure the Master Key—Every Palo Alto Networks firewall has a default master key that encrypts private keys that are used to authenticate administrators when they access management interfaces on the firewall. As a best practice to safeguard the keys, configure the master key on each firewall to be unique.
- Manage Firewall Administrators—Every Palo Alto Networks firewall and appliance is preconfigured with a default administrative account (admin) that provides full read-write access (also known as superuser access) to the device. As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This allows you to better protect the device from unauthorized configuration (or modification) and to enable logging of the actions of each individual administrator.
- Enable User Identification (User-ID)—User-ID is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses.
- Enable Decryption—Palo Alto Networks firewalls provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted or tunneled traffic.
- Enable Passive DNS Collection for Improved Threat Intelligence—Enable this opt-in feature to enable the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis in order to improve threat intelligence and threat prevention capabilities.
- Follow the Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions.

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

# Device Management

Administrators can configure, manage, and monitor the Palo Alto Networks firewalls using the web interface, the CLI, and the API management interface. Role-based administrative access to the management interfaces can be customized in order to delegate specific tasks or permissions to certain administrators. See the following topics for information on device management options, including how to begin using the management interfaces and how to customize administrator roles:

- Management Interfaces
- Manage Firewall Administrators
- Reference: Web Interface Administrator Access
- Reference: Port Numbers Used by Palo Alto Networks Devices
- Reset the Firewall to Factory Default Settings

# Management Interfaces

PAN-OS firewalls and Panorama provide three user interfaces: a web interface, a command line interface (CLI), and a XML-based management API. See the following topics for how to access and begin using each of the device management interfaces:

- Use the Web Interface to complete administrative tasks and generate reports from the web interface with relative ease. This graphical interface allows you to access the firewall using HTTPS and it is the best way to perform administrative tasks.
- Use the Command Line Interface (CLI) to type through the commands in rapid succession to complete a series of tasks. The CLI is a no-frills interface that supports two command modes and each mode has its own hierarchy of commands and statements. When you get familiar with the nesting structure and the syntax of the commands, the CLI allows quick response times and offers administrative efficiency.
- Use the XML API to streamline your operations and integrate with existing, internally developed applications and repositories. The XML API is provided as a web service that is implemented using HTTP/HTTPS requests and responses.

 © Palo Alto Networks, Inc.

# Use the Web Interface

The following topics describes how to begin using the firewall web interface. For detailed information about the tabs and fields that are available in the web interface, refer to the Web Interface Reference Guide.


- Launch the Web Interface
- Navigate the Web Interface
- Commit or Validate Changes
- Global Find
- Use Configuration Pages
- Required Fields
- Lock Transactions


## Launch the Web Interface

The following web browsers are supported for access to the web interface for PAN-OS firewalls and Panorama:

- Internet Explorer 7+
- Firefox 3.6+
- Safari 5+
- Chrome 11+

Launch an Internet browser and enter the firewall's IP address. Enter your user credentials. If logging in to the firewall for the first time, type the default **admin** into both the **Name** and **Password** fields.

To view information on how to use a specific page and an explanation of the fields and options on the page, click the **Help** icon

in the upper right area of the page to open the online help system. In addition to displaying context-sensitive help for a page, clicking the **Help** icon displays a help navigation pane with options to browse and search all help content.

## Navigate the Web Interface

The following conventions apply when using the web interface.

- To display the menu items for a general functional category, click the tab, such as **Objects** or **Device**, near the top of the browser window.

- Click an item on the side menu to display a panel.

- To display submenu items, click the ⊞ icon to the left of an item. To hide submenu items, click the ⊟ icon to the left of the item.

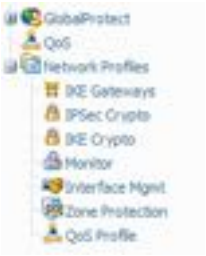


- To search the candidate configuration on a firewall or on Panorama for a particular string, click the **Search** icon to start a Global Find search.

- On most configuration pages, you can click **Add** to create a new item.

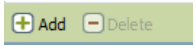


- To delete one or more items, select their check boxes and click **Delete**. In most cases, the system prompts you to confirm by clicking **OK** or to cancel the deletion by clicking **Cancel**.

- On some configuration pages, you can select the check box for an item and click **Clone** to create a new item with the same information as the selected item.

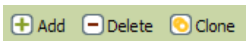


- To modify an item, click its underlined link.

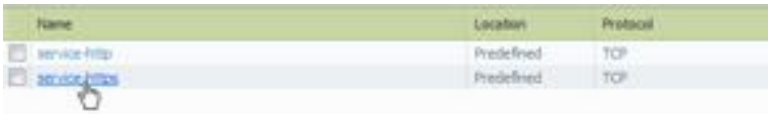


 © Palo Alto Networks, Inc.

- To view the current list of tasks, click the **Tasks** icon in the lower right corner of the page. The Task Manager window opens to show the list of tasks, along with status, start times, associated messages, and actions. Use the **Show** drop-down to filter the list of tasks.

- The web interface language is controlled by the current language of the computer that is managing the device if a specific language preference has not been defined. For example, if the computer you use to manage the firewall has a locale of Spanish, when you log in to the firewall, the web interface will be in Spanish.

- To specify a language that will always be used for a given account regardless of the locale of the computer, click the **Language** icon in the lower right corner of the page and the Language Preference window opens. Click the drop-down to select the desired language and then click **OK** to save your change.

- On pages that list information you can modify (for example, the **Setup** page on the **Devices** tab), click the icon in the upper right corner of a section to edit the settings.

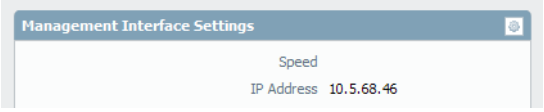


- After you configure settings, you must click **OK** or **Save** to store the changes. When you click **OK**, the current *candidate* configuration is updated.

## Commit or Validate Changes

Click **Commit** at the top of the web interface to open the Commit dialog box, which displays the following options. Certain options only display if you click **Advanced**.

If dependencies between the configuration changes you included and excluded cause a validation error, perform the commit with all the changes included. For example, if your changes introduce a new Log Forwarding profile (an object) that references a new Syslog server profile (a device setting), the commit must include both policy and object configurations and device and network configurations.

- **Include Device and Network configuration**—Include the device and network configuration changes in the commit operation.
- **Include Shared Object configuration**—(Multi-virtual system firewalls only) Include the shared object configuration changes in the commit operation.
- **Include Policy and Object configuration**—(Non-multi-virtual system firewalls only) Include the policy and object configuration changes in the commit operation.
- **Include Virtual System configuration**—Include all virtual systems or choose **Select one or more virtual systems**.
- **Preview Changes**—Click this button to display a two-pane window that shows proposed changes in the candidate configuration compared to the current running configuration. You can choose the number of lines of context to display, or show all lines. Changes are color coded based on items that you and other administrators added (green), modified (yellow), or deleted (red) since the last commit.

Because the preview results display in a new window, your browser must allow pop-ups. If the preview window does not open, refer to your browser documentation for the steps to unblock pop-ups.

- **Validate Changes**—Click this button to perform a syntactic validation (of configuration syntax) and semantic validation (whether the configuration is complete and makes sense) of a firewall or Panorama candidate configuration before committing it. The results display all of the errors and warnings of a full commit or virtual system commit, including rule shadowing and application dependency warnings. Possible errors could be an invalid route destination or a missing account and password that are required to query a server. Such validation significantly reduces failures at commit time.

## Global Find

Global Find enables you to search the candidate configuration on a firewall or on Panorama for a particular string, such as an IP address, object name, policy rule name, threat ID, or application name. The search results are grouped by category and provide links to the configuration location in the web interface, so that you can easily find all of the places where the string is referenced. The search results also help you identify other objects that depend on or make reference to the search term or string. For example, when deprecating a security profile enter the profile name in Global Find to locate all instances of the profile and then click each instance to navigate to the configuration page and make the necessary change. After all references are removed, you can then delete the profile. You can do this for any configuration item that has dependencies.

Global Find will not search dynamic content (such as logs, address ranges, or allocated DHCP addresses). In the case of DHCP, you can search on a DHCP server attribute, such as the DNS entry, but you cannot search for individual addresses allocated to users. Global Find also does not search for individual user or group names identified by User-ID unless the user/group is defined in a policy. In general, you can only search content that the firewall writes to the configuration.

 © Palo Alto Networks, Inc.

**Use Global Find**

- Launch Global Find by clicking the **Search** icon located on the upper right of the web interface.

- To access the Global Find from within a configuration area, click the drop-down next to an item and click **Global Find** as follows:

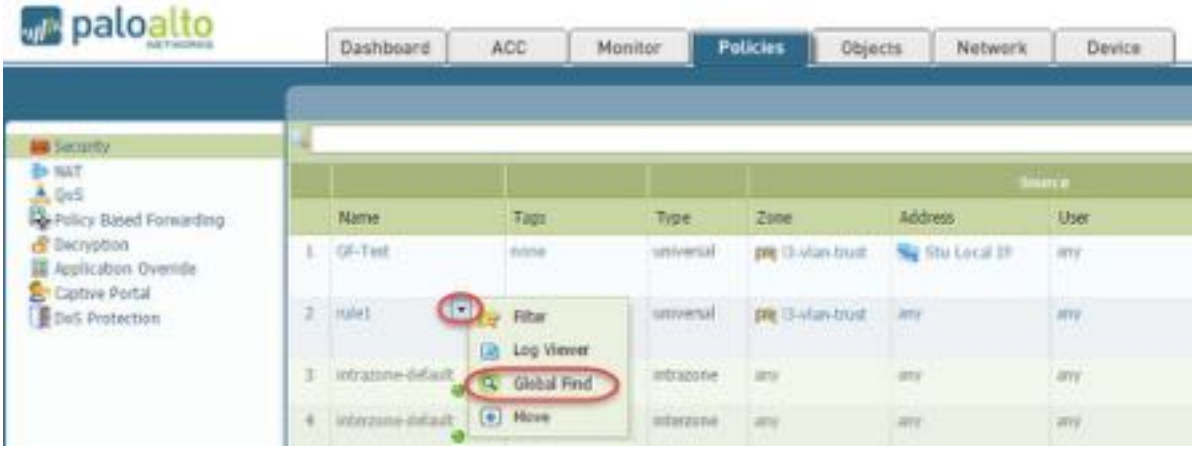

**Use Global Find (Continued)**

For example, click **Global Find** on a zone named **`l3-vlan-trust`** to search the candidate configuration for each location where the zone is referenced. The following screen capture shows the search results for the zone l3-vlan-trust:

**Search tips:**

- If you initiate a search on a firewall that has multiple virtual systems enabled or if custom Administrative Roles are defined, Global Find will only return results for areas of the firewall in which the administrator has permissions. The same applies to Panorama device groups.
- Spaces in search terms are handled as AND operations. For example, if you search on `corp policy`, the search results include instances where **corp** and **policy** exist in the configuration.
- To find an exact phrase, enclose the phrase in quotation marks.
- To rerun a previous search, click the Search icon located on the upper right of the web interface and a list of the last 20 searches will be displayed. Click an item in the list to rerun that search. The search history list is unique to each administrator account.

## Use Configuration Pages

The tables on configuration pages include sorting and column chooser options. Click a column header to sort on that column, and click again to change the sort order. Click the arrow to the right of any column and select check boxes to choose the columns to display.

## Required Fields

Required fields are shown with a light yellow background. A message indicating that the field is required appears when you hover over or click in the field entry area.

 © Palo Alto Networks, Inc.

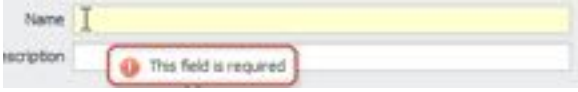


## Lock Transactions

The web interface provides support for multiple administrators by allowing an administrator to lock a current set of transactions, thereby preventing configuration changes or commit operations by another administrator until the lock is removed. The following types of locks are supported:

- **Config lock**—Blocks other administrators from making changes to the configuration. This type of lock can be set globally or for a virtual system. It can be removed only by the administrator who set it or by a superuser on the system.
- **Commit Lock**—Blocks other administrators from committing changes until all of the locks have been released. This type of lock prevents collisions that can occur when two administrators are making changes at the same time and the first administrator finishes and commits changes before the second administrator has finished. The lock is released when the current changes are committed by the administrator who applied the lock, or it can be released manually.

Any administrator can open the lock window to view the current transactions that are locked, along with a time stamp for each.

To lock a transaction, click the unlocked icon on the top bar to open the Locks dialog box. Click **Take a Lock**, select the scope of the lock from the drop-down, and click **OK**. Add additional locks as needed, and then click **Close** to close the Lock dialog box.

The transaction is locked, and the icon on the top bar changes to a locked icon that shows the number of locked items in parentheses.

To unlock a transaction, click the locked icon on the top bar to open the Locks window. Click the icon for the lock that you want to remove, and click **Yes** to confirm. Click **Close** to close the Lock dialog box.

You can arrange to automatically acquire a commit lock by selecting the **Automatically acquire commit lock** check box in the Management area of the **Device Setup** page.

# Use the Command Line Interface (CLI)

The PAN-OS CLI allows you to access Firewall and Panorama devices, view status and configuration information, and modify configurations. Access to the PAN-OS CLI is provided through SSH, Telnet, or direct console access.

The following topics describe how to access and begin using the PAN-OS CLI:

- Access the PAN-OS CLI
- Operational and Configuration Modes

## Access the PAN-OS CLI

Use a terminal emulator, such as PuTTY, to connect to the CLI in one of the following ways:

- **SSH Connection**—If you Perform Initial Configuration, you can establish a CLI connection over the network using a secure shell (SSH) connection.
- **Serial Connection**—If you have not yet completed initial configuration or if you chose not to enable SSH on the firewall, you can establish a direct serial connection from a serial interface on your management computer to the Console port on the firewall.

| **Access the PAN-OS CLI** | |
|---|---|
| Step 1 | Launch the terminal emulation software and select the type of connection (Serial or SSH).<br>• To establish an SSH connection, enter the hostname or IP address of the firewall or Panorama you want to connect to and set the port to 22.<br>• To establish a Serial connection, connect a serial interface on management computer to the Console port on the firewall. Configure the Serial connection settings in the terminal emulation software as follows:<br>  • Data rate: 9600<br>  • Data bits: 8<br>  • Parity: none<br>  • Stop bits: 1<br>  • Flow control: none |
| Step 2 | When prompted to log in, enter your administrative username.<br>The default superuser username is `admin`. To set up CLI access for other administrative users, Configure an Administrative Account. |
| Step 3 | Enter the administrative password.<br>The default superuser password is `admin`. However, for security reasons you should immediately change the admin password.<br>The CLI opens in Operational mode, and the CLI prompt is displayed:<br>username@hostname><br>You can tell you are in Operational mode because the command prompt ends with a >. |

 © Palo Alto Networks, Inc.

## Operational and Configuration Modes

When you log in, the PAN-OS CLI opens in operational mode. You can move between operational and configuration modes at any time. Use operational mode to view the state of the system, navigate the PAN-OS CLI, and enter configuration mode. Use configuration mode to view and modify the configuration hierarchy.

- To enter configuration mode from operational mode, use the `configure` command:

```
username@hostname> configure
Entering configuration mode
[edit]
username@hostname#
```

- To leave configuration mode and return to operational mode, use the `quit` or `exit` command:

```
username@hostname# quit
Exiting configuration mode
username@hostname>
```

- To enter an operational mode command while in configuration mode, use the `run` command, for example:

```
username@hostname# run ping host 10.1.1.2
PING 10.1.1.2 (10.1.1.2) 56(84) bytes of data
...
username@hostname#
```

- To direct an Operational mode command to a particular VSYS, specify the target VSYS with the following command:

```
username@hostname# set system setting target-vsys <vsys_name>
```

# Use the XML API

Palo Alto Networks XML API uses standard HTTP requests to send and receive data, allowing access to several types of data on the device so the data can be easily integrated with and used in other systems. Use the XML-based Management API to view a firewall or Panorama's configuration, extract report data in XML format, and execute operational commands. API calls can be made directly from command line utilities such as cURL or wget, or using any scripting or application framework that supports RESTful services. When using the API with command lines tools, both HTTP GET and POST methods are supported.

You must generate an API key in order to use the XML API. The API key authenticates the user to the firewall, application, or Panorama. After you have generated an API key, you can use the key to perform device configuration and operational tasks, retrieve reports and logs, and import and export files. See Generate an API Key for steps to generate an API key.

The following table shows the URL structure for API requests:

| PAN-OS Version | XML API URL Structure |
|---|---|
| Prior to PAN-OS 4.1.0 | http(s)://hostname/esp/restapi.esp?request-parameters-values |
| PAN-OS 4.1.0 and later | http(s)://hostname/api/?request-parameters-values |

URL structure item definitions:

- **hostname**—Device's IP address or Domain name.
- **request-parameters-values**—A series of multiple 'parameter=value' pairs separated by the ampersand character (&). These values can either be keywords or data-values in standard or XML format (response data is always in XML format).

There are APIs for PAN-OS, User-ID, and WildFire products. For more information on how to use the API interface, refer to the PAN-OS XML API Usage Guide.

## Generate an API Key

In order to use the API to manage a firewall or application, an API key is required to authenticate all API calls. Admin account credentials are used to generate API keys.

As a best practice, create a separate admin account for XML-based administration.

| Generate an API key | |
|---|---|
| Step 1 Create an administrator account. | 1. In the web interface, on the **Device > Administrators** tab, click **Add**.<br>2. Enter a login **Name** for the admin.<br>3. Enter and confirm a **Password** for the admin.<br>4. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

| Generate an API key (Continued) | | |
|---|---|---|
| Step 2 | Request an API key. | Replace the hostname, username and password parameters in the following URL with the appropriate values from your administrator account credentials:<br>http(s)://hostname/api/?type=keygen&user=username&password=password<br>The API key is displayed in an XML block. For example:<br><response status="success"><br><result><br><key>0RgWc42Oi0vDx2WRUIUM6A</key><br></result><br></response> |
| Step 3 | (Optional) Revoke or change an API key.<br>For PAN-OS 4.1.0 and later releases, generating an API key using the same administrator account credentials returns unique API keys every time, and all of the keys are valid.<br>You can choose to revoke and then change an API key associated with an administrator account by changing the password associated with the administrator account. Any API keys that were generated using the previous credentials would no longer be valid. | 1. On the **Device > Adminstrators** tab, open the administrator account associated with the API key.<br>2. Enter and confirm a new **Password** for the administrator account.<br>3. Click **OK** and **Commit**.<br>Any API keys associated with the admin account prior to the password change are revoked upon Commit.<br>4. (Optional) Use the updated administrator account credentials to generate a new API key. See Step 2. |

Example work flow using an API key:

Request an API key by entering the URL with the appropriate values in a web browser:

```
https://10.xx.10.50/esp/restapi.esp?type=keygen&user=admin&password=admin
```

Entering the URL displays an XML block that contains the API key:

```
<response status="success">
  <result>
    <key>0RgWc42Oi0vDx2WRUIUM6A</key>
  </result>
</response>
```

Continue to use the API key to create API requests. For example, to generate a report:

```
https://10.xx.10.50/esp/restapi.esp?type=report&reporttype=dynamic&reportname
=top-app-summary&period=last-hour&topn=5&key=0RgWc42Oi0vDx2WRUIUM6A=
```

# Manage Firewall Administrators

Administrative accounts specify roles and authentication methods for the administrators of Palo Alto Networks firewalls. Every Palo Alto Networks firewall has a predefined default administrative account (admin) that provides full read-write access (also known as superuser access) to the firewall.

As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This enables you to better protect the firewall from unauthorized configuration and enables logging of the actions of individual administrators.

- Administrative Roles
- Administrative Authentication
- Configure Administrative Accounts and Authentication

 © Palo Alto Networks, Inc.

# Administrative Roles

A *role* defines the type of access that an administrator has to the firewall.

▲ Administrative Role Types
▲ Configure an Admin Role Profile

## Administrative Role Types

The role types are:

- **Dynamic Roles**—These are built-in roles that provide access to the firewall. When new features are added, the firewall automatically updates the definitions of dynamic roles; you never need to manually update them. The following table lists the access privileges associated with dynamic roles.

| Dynamic Role | Privileges |
|---|---|
| Superuser | Full access to the firewall, including defining new administrator accounts and virtual systems. You must have superuser privileges to create an administrative user with superuser privileges. |
| Superuser (read-only) | Read-only access to the firewall. |
| Virtual system administrator | Full access to a selected virtual system (vsys) on the firewall. |
| Virtual system administrator (read-only) | Read-only access to a selected vsys on the firewall. |
| Device administrator | Full access to all firewall settings except for defining new accounts or virtual systems. |
| Device administrator (read-only) | Read-only access to all firewall settings except password profiles (no access) and administrator accounts (only the logged in account is visible). |

- **Admin Role Profiles**—Custom roles you can configure for more granular access control over the functional areas of the web interface, CLI, and XML API. For example, you can create an Admin Role profile for your operations staff that provides access to the firewall and network configuration areas of the web interface and a separate profile for your security administrators that provides access to security policy definitions, logs, and reports. On a multi-vsys firewall, you can select whether the role defines access for all virtual systems or for a specific vsys. When new features are added to the product, you must update the roles with corresponding access privileges: the firewall does not automatically add new features to custom role definitions. For details on the privileges you can configure for custom administrator roles, see Reference: Web Interface Administrator Access.

## Configure an Admin Role Profile

Admin Role profiles enable you to define granular administrative access privileges to ensure protection for sensitive company information and privacy for end users. As a best practice, create Admin Role profiles that allow administrators to access only the areas of the management interfaces that they require to perform their jobs.

© Palo Alto Networks, Inc.

| Configure an Admin Role Profile | |
|---|---|
| Step 1 | Select **Device > Admin Roles** and click **Add**. |
| Step 2 | Enter a **Name** to identify the role. |
| Step 3 | For the scope of the **Role**, select **Device** or **Virtual System**. |
| Step 4 | In the **Web UI** and **XML API** tabs, click the icon for each functional area to toggle it to the desired setting: Enable, Read Only, or Disable. For details on the **Web UI** options, see Web Interface Access Privileges. |
| Step 5 | Select the **Command Line** tab and select a CLI access option. The **Role** scope controls the available options:<br>• **Device** role—**superuser**, **superreader**, **deviceadmin**, **devicereader**, or **None**<br>• **Virtual System** role—**vsysadmin**, **vsysreader**, or **None** |
| Step 6 | Click **OK** to save the profile. |
| Step 7 | Assign the role to an administrator. See Configure an Administrative Account. |


© Palo Alto Networks, Inc.

## Administrative Authentication

You can configure the following types of administrator authentication:

| Account Type | Authentication Method | Description |
|---|---|---|
| Local | Local database | Both the administrator account credentials and the authentication mechanisms are local to the firewall. You can further secure a local administrator account by creating a password profile that defines a validity period for passwords and by setting firewall-wide password complexity settings. If your network supports Kerberos single sign-on (SSO), you can configure local authentication as a fallback in case SSO fails. For details, see Configure Kerberos SSO and External or Local Authentication for Administrators. |
| Local | SSL-based | The administrator accounts are local to the firewall, but authentication is based on SSH certificates (for CLI access) or client certificates (for web interface access). For details, see Configure SSH Key-Based Administrator Authentication to the CLI and Configure Certificate-Based Administrator Authentication to the Web Interface. |
| Local | External service | The administrator accounts are local to the firewall, but external services (LDAP, Kerberos, TACACS+, or RADIUS) handle the authentication functions. If your network supports Kerberos single sign-on (SSO), you can configure external authentication as a fallback in case SSO fails. For details, see Configure Kerberos SSO and External or Local Authentication for Administrators. |
| External | External service | An external RADIUS server handles account management and authentication. You must define Vendor-Specific Attributes (VSAs) on your RADIUS server that map to the administrator role, access domain, user group (if applicable), and virtual system (if applicable). For details, see Configure RADIUS Vendor-Specific Attributes for Administrator Authentication. |

## Configure Administrative Accounts and Authentication

- Configure an Administrative Account
- Configure Kerberos SSO and External or Local Authentication for Administrators
- Configure Certificate-Based Administrator Authentication to the Web Interface
- Configure SSH Key-Based Administrator Authentication to the CLI
- Configure RADIUS Vendor-Specific Attributes for Administrator Authentication


© Palo Alto Networks, Inc.

## Configure an Administrative Account

Administrative accounts specify roles and authentication methods for the administrators of Palo Alto Networks firewalls.

| Configure an Administrative Account | |
|---|---|
| Step 1 | Select **Device > Administrators** and click **Add**. |
| Step 2 | Enter a user **Name**. |
| Step 3 | Select an **Authentication Profile** or sequence if you configured either for the user. The default option **None** specifies that the firewall will locally manage and authenticate the account without a local database; you must enter and confirm the **Password**. |
| Step 4 | Select the **Administrator Type**. If you configured a custom role for the user, select **Role Based** and select the Admin Role **Profile**. Otherwise, select **Dynamic** (the default) and select a dynamic role. |
| Step 5 | Click **OK** and **Commit**. |

# Configure Kerberos SSO and External or Local Authentication for Administrators

You can configure the firewall to first try Kerberos single sign-on (SSO) authentication and, if that fails, fall back to External service or Local database authentication.

| Configure Kerberos SSO and External or Local Authentication for Administrators | | |
|---|---|---|
| Step 1 | Configure a Kerberos keytab for the firewall.<br><br>Required for Kerberos SSO authentication. | Create a Kerberos keytab. A keytab is a file that contains Kerberos account information (principal name and hashed password) for the firewall. |
| Step 2 | Configure an Admin Role profile.<br><br>Required if you are assigning a custom role to the administrator. | Configure an Admin Role Profile. |
| Step 3 | Create the local database.<br><br>Required for local database authentication. | 1. Add the user account:<br>a. Select **Device > Local User Database > Users** and click **Add**.<br>b. Enter a user **Name** for the administrator.<br>c. Enter a **Password** and **Confirm Password**.<br>d. Be sure the **Enable** check box is selected and click **OK**.<br>2. (Optional) If the user is a member of a group, assign the user to that group:<br>a. Select **Device > Local User Database > User Groups** and click the Name of an existing group to edit it, or click **Add** to create a new group.<br>b. Enter a **Name** to identify the group.<br>c. Click **Add**, select the user you just created, and click **OK**. |
| Step 4 | Configure access to an external authentication service.<br><br>Required for external authentication. | Configure a server profile for the authentication service type:<br>• Configure a RADIUS Server Profile.<br>• Configure a TACACS+ Server Profile.<br>• Configure an LDAP Server Profile.<br>• Configure a Kerberos Server Profile. |
| Step 5 | Configure an authentication profile.<br><br>If your users are in multiple Kerberos realms, create an authentication profile for each realm and assign all the profiles to an authentication sequence. You can then assign the same authentication sequence to all user accounts (Step 6). | Configure an Authentication Profile and Sequence. |

 © Palo Alto Networks, Inc.

| Configure Kerberos SSO and External or Local Authentication for Administrators (Continued) | |
|---|---|
| Step 6 Configure an administrator account. | Configure an Administrative Account.<br>• For local database authentication, specify the **Name** of the user you defined in Step 3.<br>• Assign the **Authentication Profile** or sequence and the Admin Role **Profile** that you just created. |

# Configure Certificate-Based Administrator Authentication to the Web Interface

As a more secure alternative to password-based authentication to the web interface of a Palo Alto Networks firewall, you can configure certificate-based authentication for administrator accounts that are local to the firewall. Certificate-based authentication involves the exchange and verification of a digital signature instead of a password.

Configuring certificate-based authentication for any administrator disables the username/password logins for all administrators on the firewall; administrators thereafter require the certificate to log in.

| Configure Certificate-Based Administrator Authentication to the Web Interface | | |
|---|---|---|
| Step 1 | Generate a certificate authority (CA) certificate on the firewall.<br>You will use this CA certificate to sign the client certificate of each administrator. | Create a Self-Signed Root CA Certificate.<br>Alternatively, you can Import a Certificate and Private Key from your enterprise CA. |
| Step 2 | Configure a certificate profile for securing access to the web interface. | Configure a Certificate Profile.<br>• Set the **Username Field** to **Subject**.<br>• Select **Add** in the CA Certificates section and select the **CA Certificate** you just created or imported. |
| Step 3 | Configure the firewall to use the certificate profile for authenticating administrators. | 1. Select **Device > Setup > Management** and edit the Authentication Settings.<br>2. Select the **Certificate Profile** you just created and click **OK**. |
| Step 4 | Configure the administrator accounts to use client certificate authentication. | For each administrator who will access the firewall web interface, Configure an Administrative Account. Select the **Use only client certificate authentication** check box.<br>If you have already deployed client certificates that your enterprise CA generated, skip to Step 8. Otherwise, go to Step 5. |
| Step 5 | Generate a client certificate for each administrator. | Generate a Certificate on the Device. In the **Signed By** drop-down, select the CA certificate you created. |
| Step 6 | Export the client certificate. | 1. Export a Certificate and Private Key.<br>2. **Commit** your changes. The firewall restarts and terminates your login session. Thereafter, administrators can access the web interface only from client systems that have the client certificate you generated. |
| Step 7 | Import the client certificate into the client system of each administrator who will access the web interface. | Refer to your web browser documentation. |


© Palo Alto Networks, Inc.

| Configure Certificate-Based Administrator Authentication to the Web Interface (Continued) | | |
|---|---|---|
| Step 8 | Verify that administrators can access the web interface. | 1. Open the firewall IP address in a browser on the computer that has the client certificate.<br>2. When prompted, select the certificate you imported and click **OK**. The browser displays a certificate warning.<br>3. Add the certificate to the browser exception list.<br>4. Click **Login**. The web interface should appear without prompting you for a username or password. |

## Configure SSH Key-Based Administrator Authentication to the CLI

For administrators who use Secure Shell (SSH) to access the CLI of a Palo Alto Networks firewall, SSH keys provide a more secure authentication method than passwords. SSH keys almost eliminate the risk of brute-force attacks, provide the option for two-factor authentication (private key and passphrase), and don't send passwords over the network. SSH keys also enable automated scripts to access the CLI.

| Configure SSH Key-Based Administrator Authentication to the CLI | | |
|---|---|---|
| Step 1 | Use an SSH key generation tool to create an asymmetric keypair on the client system of the administrator.<br>The supported key formats are IETF SECSH and Open SSH. The supported algorithms are DSA (1,024 bits) and RSA (768-4,096 bits). | For the commands to generate the keypair, refer to your SSH client documentation.<br>The public key and private key are separate files. Save both to a location that the firewall can access. For added security, enter a passphrase to encrypt the private key. The firewall prompts the administrator for this passphrase during login. |
| Step 2 | Configure the administrator account to use public key authentication. | 1. Configure an Administrative Account.<br>• Configure the authentication method to use as a fallback if SSH key authentication fails. If you configured an **Authentication Profile** for the administrator, select it in the drop-down. If you select **None**, you must enter a **Password** and **Confirm Password**.<br>• Select the **Use Public Key Authentication (SSH)** check box, click **Import Key**, **Browse** to the public key you just generated, and click **OK**.<br>2. **Commit** your changes. |
| Step 3 | Configure the SSH client to use the private key to authenticate to the firewall. | Perform this task on the client system of the administrator. For the steps, refer to your SSH client documentation. |

| Configure SSH Key-Based Administrator Authentication to the CLI (Continued) | |
|---|---|
| Step 4 Verify that the administrator can access the firewall CLI using SSH key authentication. | 1. Use a browser on the client system of the administrator to go to the firewall IP address.<br>2. Log in to the firewall CLI as the administrator. After entering a username, you will see the following output (the key value is an example):<br>`Authenticating with public key "dsa-key-20130415"`<br>3. If prompted, enter the passphrase you defined when creating the keys. |

## Configure RADIUS Vendor-Specific Attributes for Administrator Authentication

The following procedure provides an overview of the tasks required to use RADIUS Vendor-Specific Attributes (VSAs) for administrator authentication to Palo Alto Networks firewalls. For detailed instructions, refer to the following Knowledge Base articles:

- For Windows 2003 Server, Windows 2008 (and later), and Cisco ACS 4.0—RADIUS Vendor-Specific Attributes (VSAs)
- For Cisco ACS 5.2—Configuring Cisco ACS 5.2 for use with Palo Alto VSA

Before starting this procedure, you must:

- ❒ Create the administrative accounts in the directory service that your network uses (for example, Active Directory).
- ❒ Set up a RADIUS server that can communicate with that directory service.

 © Palo Alto Networks, Inc.

| Use RADIUS Vendor-Specific Attributes for Account Authentication | |
|---|---|
| Step 1 Configure the firewall. | 1. Configure an Admin Role Profile if the administrator will use a custom role.<br>2. Configure an access domain if the firewall has more than one virtual system (vsys):<br>a. Select **Device > Access Domain**, **Add** and access domain and enter a **Name** to identify it.<br>b. **Add** each vsys that the administrator will access, and then click **OK**.<br>3. Configure a RADIUS Server Profile.<br>4. Configure an authentication profile. Set the authentication **Type** to **RADIUS** and assign the RADIUS **Server Profile**.<br>5. Configure the firewall to use the authentication profile for administrator access—Select **Device > Setup > Management**, edit the Authentication Settings, and select the **Authentication Profile**.<br>6. Click **OK** and **Commit**. |
| Step 2 Configure the RADIUS server. | 1. Add the firewall IP address or hostname as the RADIUS client.<br>2. Define the VSAs for administrator authentication. You must specify the vendor code (25461 for Palo Alto Networks firewalls) and the VSA name, number, and value: see RADIUS Vendor-Specific Attributes for Palo Alto Networks Devices. |

# Reference: Web Interface Administrator Access

You can configure privileges for an entire device or for one or more virtual systems (on platforms that support multiple virtual systems). Within that **Device** or **Virtual System** designation, you can configure privileges for custom administrator roles, which are more granular than the fixed privileges associated with a dynamic administrator role.

Configuring privileges at a granular level ensures that lower level administrators cannot access certain information. You can create custom roles for firewall administrators (see Configure an Administrative Account), Panorama administrators, or Device Group and Template administrators (refer to the Panorama Administrator's Guide). The following topics describe the privileges you can configure for custom administrator roles.

▲ Web Interface Access Privileges

▲ Panorama Web Interface Access

## Web Interface Access Privileges

If you want to prevent a role-based administrator from accessing specific tabs on the web interface, you can disable the tab and the administrator will not even see it when logging in using the associated role-based administrative account. For example, you could create an Admin Role Profile for your operations staff that provides access to the **Device** and **Network** tabs only and a separate profile for your security administrators that provides access to the **Object**, **Policy**, and **Monitor** tabs.

An admin role can apply at the **Device** level or **Virtual System** level; the choice is made in the Admin Role Profile by clicking the **Device** or **Virtual System** radio button. If the **Virtual System** button is selected, the admin assigned this profile is restricted to the virtual system(s) he or she is assigned to. Furthermore, only the **Device > Setup > Services > Virtual Systems** tab is available to that admin, not the **Global** tab.

The following table describes the tab-level access privileges you can assign to the admin role profile at the **Device** level. It also provides cross-references to additional tables that detail granular privileges within a tab.

You can also configure an admin role profile to:

- Define User Privacy Settings in the Admin Role Profile
- Restrict Admin Access to Commit Functions
- Restrict Admin Access to Validate Functions
- Provide Granular Access to Global Settings

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Dashboard | Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| ACC | Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Yes | No | Yes |
| Monitor | Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the admin can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab. | Yes | No | Yes |
| Policies | Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the admin can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab. | Yes | No | Yes |
| Objects | Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab. | Yes | No | Yes |
| Network | Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the admin can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab. | Yes | No | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Device | Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any device-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab.<br>You cannot enable access to the **Admin Roles** or **Administrators** nodes for a role-based administrator even if you enable full access to the **Device** tab. | Yes | No | Yes |

## Provide Granular Access to the Monitor Tab

In some cases you might want to enable the administrator to view some but not all areas of the **Monitor** tab. For example, you might want to restrict operations administrators to the Config and System logs only, because they do not contain sensitive user data. Although this section of the admin role definition specifies what areas of the **Monitor** tab the administrator can see, you can also couple privileges in this section with privacy privileges, such as disabling the ability to see usernames in logs and reports. One thing to keep in mind, however, is that any system-generated reports will still show usernames and IP addresses even if you disable that functionality in the role. For this reason, if you do not want the administrator to see any of the private user information, disable access to the specific reports as detailed in the following table.

The following table lists the **Monitor** tab access levels and the administrator roles for which they are available.

Device Group and Template roles can see log data only for the device groups that are within the access domains assigned to those roles.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Monitor | Enables or disables access to the **Monitor** tab. If disabled, the admin will not see this tab or any of the associated logs or reports. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Logs | Enables or disables access to all log files. You can also leave this privilege enabled and then disable specific logs that you do not want the admin to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the logs, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Traffic | Specifies whether the admin can see the traffic logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Threat | Specifies whether the admin can see the threat logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| URL Filtering | Specifies whether the admin can see the URL filtering logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| WildFire Submissions | Specifies whether the admin can see the WildFire logs. These logs are only available if you have a WildFire subscription. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Data Filtering | Specifies whether the admin can see the data filtering logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| HIP Match | Specifies whether the admin can see the HIP Match logs. HIP Match logs are only available if you have a GlobalProtect portal license and gateway subscription. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Configuration | Specifies whether the admin can see the configuration logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| System | Specifies whether the admin can see the system logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |

© Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Alarms | Specifies whether the admin can see system-generated alarms. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Automated Correlation Engine | Enables or disables access to the correlation objects and correlated event logs generated on the firewall. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Correlation Objects | Specifies whether the admin can view and enable/disable the correlation objects. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Correlated Events | Specifies whether the admin | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Packet Capture | Specifies whether the admin can see packet captures (pcaps) from the **Monitor** tab. Keep in mind that packet captures are raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the pcap and you should therefore disable the Packet Capture privilege if you are concerned about user privacy. | Firewall: Yes<br>Panorama: No<br>Device Group/Template: No | Yes | Yes | Yes |
| App Scope | Specifies whether the admin can see the App Scope visibility and analysis tools. Enabling App Scope enables access to all of the **App Scope** charts. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Session Browser | Specifies whether the admin can browse and filter current running sessions on the firewall. Keep in mind that the session browser shows raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the session browser and you should therefore disable the **Session Browser** privilege if you are concerned about user privacy. | Firewall: Yes<br>Panorama: No<br>Device Group/Template: No | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Botnet | Specifies whether the admin can generate and view botnet analysis reports or view botnet reports in read-only mode. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in scheduled botnet reports and you should therefore disable the **Botnet** privilege if you are concerned about user privacy. | Firewall: Yes<br>Panorama: No<br>Device Group/Template: No | Yes | Yes | Yes |
| PDF Reports | Enables or disables access to all PDF reports. You can also leave this privilege enabled and then disable specific PDF reports that you do not want the admin to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Manage PDF Summary | Specifies whether the admin can view, add or delete PDF summary report definitions. With read-only access, the admin can see PDF summary report definitions, but not add or delete them. If you disable this option, the admin can neither view the report definitions nor add/delete them. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |
| PDF Summary Reports | Specifies whether the admin can see the generated PDF Summary reports in **Monitor > Reports**. If you disable this option, the **PDF Summary Reports** category will not display in the **Reports** node. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| User Activity Report | Specifies whether the admin can view, add or delete User Activity report definitions and download the reports. With read-only access, the admin can see User Activity report definitions, but not add, delete, or download them. If you disable this option, the admin cannot see this category of PDF report. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |
| Report Groups | Specifies whether the admin can view, add or delete report group definitions. With read-only access, the admin can see report group definitions, but not add or delete them. If you disable this option, the admin cannot see this category of PDF report. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Email Scheduler | Specifies whether the admin can schedule report groups for email. Because the generated reports that get emailed may contain sensitive user data that is not removed by disabling the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** options and because they may also show log data to which the admin does not have access, you should disable the **Email Scheduler** option if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |
| Manage Custom Reports | Enables or disables access to all custom report functionality. You can also leave this privilege enabled and then disable specific custom report categories that you do not want the admin to be able to access. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option.<br> Reports that are scheduled to run rather than run on demand will show IP address and user information. In this case, be sure to restrict access to the corresponding report areas. In addition, the custom report feature does not restrict the ability to generate reports that contain log data contained in logs that are excluded from the admin role. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Application Statistics | Specifies whether the admin can create a custom report that includes data from the application statistics database. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Data Filtering Log | Specifies whether the admin can create a custom report that includes data from the Data Filtering logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Threat Log | Specifies whether the admin can create a custom report that includes data from the Threat logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Threat Summary | Specifies whether the admin can create a custom report that includes data from the Threat Summary database. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Traffic Log | Specifies whether the admin can create a custom report that includes data from the Traffic logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Traffic Summary | Specifies whether the admin can create a custom report that includes data from the Traffic Summary database. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| URL Log | Specifies whether the admin can create a custom report that includes data from the URL Filtering logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Hipmatch | Specifies whether the admin can create a custom report that includes data from the HIP Match logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| WildFire Log | Specifies whether the admin can create a custom report that includes data from the WildFire logs. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| View Scheduled Custom Reports | Specifies whether the admin can view a custom report that has been scheduled to generate. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| View Predefined Application Reports | Specifies whether the admin can view Application Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| View Predefined Threat Reports | Specifies whether the admin can view Threat Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| View Predefined URL Filtering Reports | Specifies whether the admin can view URL Filtering Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| View Predefined Traffic Reports | Specifies whether the admin can view Traffic Reports. Privacy privileges do not impact reports available on the **Monitor > Reports** node and you should therefore disable access to the reports if you have user privacy requirements. | Firewall: Yes<br>Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |

## Provide Granular Access to the Policy Tab

If you enable the Policy option in the admin role profile, you can then enable, disable, or provide read-only access to specific nodes within the tab as necessary for the admin role you are defining. By enabling access to a specific policy type, you enable the ability to view, add, or delete policy rules. By enabling read-only access to a specific policy, you enable the admin to view the corresponding policy rule base, but not add or delete rules. Disabling access to a specific type of policy prevents the admin from seeing the policy rule base.

Because policy that is based on specific users (by user name or IP address) must be explicitly defined, privacy settings that disable the ability to see full IP addresses or user names do not apply to the Policy tab. Therefore, you should only allow access to the Policy tab to administrators that are excluded from user privacy restrictions.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Security | Enable this privilege to allow the admin to view, add, and/or delete security rules. Set the privilege to read-only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the security rulebase, disable this privilege. | Yes | Yes | Yes |
| NAT | Enable this privilege to allow the admin to view, add, and/or delete NAT rules. Set the privilege to read-only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the NAT rulebase, disable this privilege. | Yes | Yes | Yes |
| QoS | Enable this privilege to allow the admin to view, add, and/or delete QoS rules. Set the privilege to read-only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the QoS rulebase, disable this privilege. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Policy Based Forwarding | Enable this privilege to allow the admin to view, add, and/or delete Policy-Based Forwarding (PBF) rules. Set the privilege to read-only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the PBF rulebase, disable this privilege. | Yes | Yes | Yes |
| Decryption | Enable this privilege to allow the admin to view, add, and/or delete decryption rules. Set the privilege to read-only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the decryption rulebase, disable this privilege. | Yes | Yes | Yes |
| Application Override | Enable this privilege to allow the admin to view, add, and/or delete application override policy rules. Set the privilege to read-only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the application override rulebase, disable this privilege. | Yes | Yes | Yes |
| Captive Portal | Enable this privilege to allow the admin to view, add, and/or delete Captive Portal rules. Set the privilege to read-only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the Captive Portal rulebase, disable this privilege. | Yes | Yes | Yes |
| DoS Protection | Enable this privilege to allow the admin to view, add, and/or delete DoS protection rules. Set the privilege to read-only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the DoS protection rulebase, disable this privilege. | Yes | Yes | Yes |

## Provide Granular Access to the Objects Tab

An *object* is a container that groups specific policy filter values—such as IP addresses, URLs, applications, or services—for simplified rule definition. For example, an address object might contain specific IP address definitions for the web and application servers in your DMZ zone.

When deciding whether to allow access to the objects tab as a whole, determine whether the admin will have policy definition responsibilities. If not, the admin probably does not need access to the tab. If, however, the admin will need to create policy, you can enable access to the tab and then provide granular access privileges at the node level.

By enabling access to a specific node, you give the admin the privilege to view, add, and delete the corresponding object type. Giving read-only access allows the admin to view the already defined objects, but not create or delete any. Disabling a node prevents the admin from seeing the node in the web interface.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Addresses | Specifies whether the admin can view, add, or delete address objects for use in security policy. | Yes | Yes | Yes |
| Address Groups | Specifies whether the admin can view, add, or delete address group objects for use in security policy. | Yes | Yes | Yes |
| Regions | Specifies whether the admin can view, add, or delete regions objects for use in security, decryption, or DoS policy. | Yes | Yes | Yes |
| Applications | Specifies whether the admin can view, add, or delete application objects for use in policy. | Yes | Yes | Yes |
| Application Groups | Specifies whether the admin can view, add, or delete application group objects for use in policy. | Yes | Yes | Yes |
| Application Filters | Specifies whether the admin can view, add, or delete application filters for simplification of repeated searches. | Yes | Yes | Yes |
| Services | Specifies whether the admin can view, add, or delete service objects for use in creating policy rules that limit the port numbers an application can use. | Yes | Yes | Yes |
| Service Groups | Specifies whether the admin can view, add, or delete service group objects for use in security policy. | Yes | Yes | Yes |
| Tags | Specifies whether the admin can view, add, or delete tags that have been defined on the device. | Yes | Yes | Yes |
| GlobalProtect | Specifies whether the admin can view, add, or delete HIP objects and profiles. You can restrict access to both types of objects at the GlobalProtect level, or provide more granular control by enabling the GlobalProtect privilege and restricting HIP Object or HIP Profile access. | Yes | No | Yes |
| HIP Objects | Specifies whether the admin can view, add, or delete HIP objects, which are used to define HIP profiles. HIP Objects also generate HIP Match logs. | Yes | Yes | Yes |
| HIP Profiles | Specifies whether the admin can view, add, or delete HIP Profiles for use in security policy and/or for generating HIP Match logs. | Yes | Yes | Yes |
| Dynamic Block Lists | Specifies whether the admin can view, add, or delete dynamic block lists for use in security policy. | Yes | Yes | Yes |
| Custom Objects | Specifies whether the admin can see the custom spyware and vulnerability signatures. You can restrict access to either enable or disable access to all custom signatures at this level, or provide more granular control by enabling the Custom Objects privilege and then restricting access to each type of signature. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Data Patterns | Specifies whether the admin can view, add, or delete custom data pattern signatures for use in creating custom Vulnerability Protection profiles. | Yes | Yes | Yes |
| Spyware | Specifies whether the admin can view, add, or delete custom spyware signatures for use in creating custom Vulnerability Protection profiles. | Yes | Yes | Yes |
| Vulnerability | Specifies whether the admin can view, add, or delete custom vulnerability signatures for use in creating custom Vulnerability Protection profiles. | Yes | Yes | Yes |
| URL Category | Specifies whether the admin can view, add, or delete custom URL categories for use in policy. | Yes | Yes | Yes |
| Security Profiles | Specifies whether the admin can see security profiles. You can restrict access to either enable or disable access to all security profiles at this level, or provide more granular control by enabling the Security Profiles privilege and then restricting access to each type of profile. | Yes | No | Yes |
| Antivirus | Specifies whether the admin can view, add, or delete antivirus profiles. | Yes | Yes | Yes |
| Anti-Spyware | Specifies whether the admin can view, add, or delete Anti-Spyware profiles. | Yes | Yes | Yes |
| Vulnerability Protection | Specifies whether the admin can view, add, or delete Vulnerability Protection profiles. | Yes | Yes | Yes |
| URL Filtering | Specifies whether the admin can view, add, or delete URL filtering profiles. | Yes | Yes | Yes |
| File Blocking | Specifies whether the admin can view, add, or delete file blocking profiles. | Yes | Yes | Yes |
| Data Filtering | Specifies whether the admin can view, add, or delete data filtering profiles. | Yes | Yes | Yes |
| DoS Protection | Specifies whether the admin can view, add, or delete DoS protection profiles. | Yes | Yes | Yes |
| Security Profile Groups | Specifies whether the admin can view, add, or delete security profile groups. | Yes | Yes | Yes |
| Log Forwarding | Specifies whether the admin can view, add, or delete log forwarding profiles. | Yes | Yes | Yes |
| Decryption Profile | Specifies whether the admin can view, add, or delete decryption profiles. | Yes | Yes | Yes |
| Schedules | Specifies whether the admin can view, add, or delete schedules for limiting a security policy to a specific date and/or time range. | Yes | Yes | Yes |

© Palo Alto Networks, Inc.

## Provide Granular Access to the Network Tab

When deciding whether to allow access to the **Network** tab as a whole, determine whether the admin will have network administration responsibilities, including GlobalProtect administration. If not, the admin probably does not need access to the tab.

You can also define access to the **Network** tab at the node level. By enabling access to a specific node, you give the admin the privilege to view, add, and delete the corresponding network configurations. Giving read-only access allows the admin to view the already-defined configuration, but not create or delete any. Disabling a node prevents the admin from seeing the node in the web interface.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Interfaces | Specifies whether the admin can view, add, or delete interface configurations. | Yes | Yes | Yes |
| Zones | Specifies whether the admin can view, add, or delete zones. | Yes | Yes | Yes |
| VLANs | Specifies whether the admin can view, add, or delete VLANs. | Yes | Yes | Yes |
| Virtual Wires | Specifies whether the admin can view, add, or delete virtual wires. | Yes | Yes | Yes |
| Virtual Routers | Specifies whether the admin can view, add, modify or delete virtual routers. | Yes | Yes | Yes |
| IPSec Tunnels | Specifies whether the admin can view, add, modify, or delete IPSec Tunnel configurations. | Yes | Yes | Yes |
| DHCP | Specifies whether the admin can view, add, modify, or delete DHCP server and DHCP relay configurations. | Yes | Yes | Yes |
| DNS Proxy | Specifies whether the admin can view, add, modify, or delete DNS proxy configurations. | Yes | Yes | Yes |
| GlobalProtect | Specifies whether the admin can view, add, modify GlobalProtect portal and gateway configurations. You can disable access to the GlobalProtect functions entirely, or you can enable the GlobalProtect privilege and then restrict the role to either the portal or gateway configuration areas. | Yes | No | Yes |
| Portals | Specifies whether the admin can view, add, modify, or delete GlobalProtect portal configurations. | Yes | Yes | Yes |
| Gateways | Specifies whether the admin can view, add, modify, or delete GlobalProtect gateway configurations. | Yes | Yes | Yes |
| MDM | Specifies whether the admin can view add, modify, or delete GlobalProtect MDM server configurations. | Yes | Yes | Yes |
| QoS | Specifies whether the admin can view add, modify, or delete QoS configurations. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| LLDP | Specifies whether the admin can view add, modify, or delete LLDP configurations. | Yes | Yes | Yes |
| Network Profiles | Sets the default state to enable or disable for all of the Network settings described below. | Yes | No | Yes |
| IKE Gateways | Controls access to the **Network Profiles > IKE Gateways** node. If you disable this privilege, the administrator will not see the **IKE Gateways** node or define gateways that include the configuration information necessary to perform IKE protocol negotiation with peer gateway.<br>If the privilege state is set to read-only, you can view the currently configured IKE Gateways but cannot add or edit gateways. | Yes | Yes | Yes |
| GlobalProtect IPSec Crypto | Controls access to the **Network Profiles > GlobalProtect IPSec Crypto** node.<br>If you disable this privilege, the administrator will not see that node, or configure algorithms for authentication and encryption in VPN tunnels between a GlobalProtect gateway and clients.<br>If you set the privilege to read-only, the administrator can view existing GlobalProtect IPSec Crypto profiles but cannot add or edit them. | Yes | Yes | Yes |
| IPSec Crypto | Controls access to the **Network Profiles > IPSec Crypto** node. If you disable this privilege, the administrator will not see the **Network Profiles > IPSec Crypto** node or specify protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPSec SA negotiation.<br>If the privilege state is set to read-only, you can view the currently configured IPSec Crypto configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| IKE Crypto | Controls how devices exchange information to ensure secure communication. Specify the protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPsec SA negotiation (IKEv1 Phase-1). | Yes | Yes | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Monitor | Controls access to the **Network Profiles > Monitor** node. If you disable this privilege, the administrator will not see the **Network Profiles > Monitor** node or be able to create or edit a monitor profile that is used to monitor IPSec tunnels and monitor a next-hop device for policy-based forwarding (PBF) rules.<br>If the privilege state is set to read-only, you can view the currently configured monitor profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| Interface Mgmt | Controls access to the **Network Profiles > Interface Mgmt** node. If you disable this privilege, the administrator will not see the **Network Profiles > Interface Mgmt** node or be able to specify the protocols that are used to manage the firewall.<br>If the privilege state is set to read-only, you can view the currently configured Interface management profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| Zone Protection | Controls access to the **Network Profiles > Zone Protection** node. If you disable this privilege, the administrator will not see the **Network Profiles > Zone Protection** node or be able to configure a profile that determines how the firewall responds to attacks from specified security zones.<br>If the privilege state is set to read-only, you can view the currently configured Zone Protection profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| QoS Profile | Controls access to the **Network Profiles > QoS** node. If you disable this privilege, the administrator will not see the **Network Profiles > QoS** node or be able to configure a QoS profile that determines how QoS traffic classes are treated.<br>If the privilege state is set to read-only, you can view the currently configured QoS profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |
| LLDP Profile | Controls access to the **Network Profiles > LLDP** node. If you disable this privilege, the administrator will not see the **Network Profiles > LLDP** node or be able to configure an LLDP profile that controls whether the interfaces on the firewall can participate in the Link Layer Discovery Protocol.<br>If the privilege state is set to read-only, you can view the currently configured LLDP profile configuration but cannot add or edit a configuration. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

## Provide Granular Access to the Device Tab

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Setup | Controls access to the **Setup** node. If you disable this privilege, the administrator will not see the **Setup** node or have access to device-wide setup configuration information, such as Management, Operations, Service, Content-ID, Wildfire or Session setup information.<br>If the privilege state is set to read-only, you can view the current configuration but cannot make any changes. | Yes | Yes | Yes |
| Config Audit | Controls access to the **Config Audit** node. If you disable this privilege, the administrator will not see the **Config Audit** node or have access to any device-wide configuration information. | Yes | No | Yes |
| Admin Roles | Controls access to the **Admin Roles** node. This function can only be allowed for read-only access.<br>If you disable this privilege, the administrator will not see the **Admin Roles** node or have access to any device-wide information concerning admin roles configuration.<br>If you set this privilege to read-only, you can view the configuration information for all admin roles configured on the device. | No | Yes | Yes |
| Administrators | Controls access to the **Administrators** node. This function can only be allowed for read-only access.<br>If you disable this privilege, the administrator will not see the **Administrators** node or have access to information about their own admin account.<br>If you set this privilege to read-only, the administrator can view the configuration information for their own admin account. They will not see any information about other admin accounts configured on the device. | No | Yes | Yes |
| Virtual Systems | Controls access to the **Virtual Systems** node. If you disable this privilege, the administrator will not see or be able to configure virtual systems.<br>If the privilege state is set to read-only, you can view the currently configured virtual systems but cannot add or edit a configuration. | Yes | Yes | Yes |

© Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Shared Gateways | Controls access to the **Shared Gateways** node. Shared gateways allow virtual systems to share a common interface for external communications.<br>If you disable this privilege, the administrator will not see or be able to configure shared gateways.<br>If the privilege state is set to read-only, you can view the currently configured shared gateways but cannot add or edit a configuration. | Yes | Yes | Yes |
| User Identification | Controls access to the **User Identification** node. If you disable this privilege, the administrator will not see the **User Identification** node or have access to device-wide User Identification configuration information, such as User Mapping, User-ID Agents, Service, Terminal Services Agents, Group Mappings Settings or Captive Portal Settings.<br>If you set this privilege to read-only, the administrator can view configuration information for the device but is not allowed to perform any configuration procedures. | Yes | Yes | Yes |
| VM Information Source | Controls access to the **VM Information Source** node that allows you to configure the firewall/Windows User-ID agent to collect VM inventory automatically. If you disable this privilege, the administrator will not see the **VM Information Source** node.<br>If you set this privilege to read-only, the administrator can view the VM information sources configured but cannot add, edit, or delete any sources.<br>This privilege is not available to Device Group and Template administrators. | Yes | Yes | Yes |
| High Availability | Controls access to the **High Availability** node. If you disable this privilege, the administrator will not see the **High Availability** node or have access to device-wide high availability configuration information such as General setup information or Link and Path Monitoring.<br>If you set this privilege to read-only, the administrator can view High Availability configuration information for the device but is not allowed to perform any configuration procedures. | Yes | Yes | Yes |
| Certificate Management | Sets the default state to enable or disable for all of the Certificate settings described below. | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Certificates | Controls access to the **Certificates** node. If you disable this privilege, the administrator will not see the **Certificates** node or be able to configure or access information regarding Device Certificates or Default Trusted Certificate Authorities.<br>If you set this privilege to read-only, the administrator can view Certificate configuration information for the device but is not allowed to perform any configuration procedures. | Yes | Yes | Yes |
| Certificate Profile | Controls access to the **Certificate Profile** node. If you disable this privilege, the administrator will not see the **Certificate Profile** node or be able to create certificate profiles.<br>If you set this privilege to read-only, the administrator can view Certificate Profiles that are currently configured for the device but is not allowed to create or edit a certificate profile. | Yes | Yes | Yes |
| OCSP Responder | Controls access to the **OCSP Responder** node. If you disable this privilege, the administrator will not see the **OCSP Responder** node or be able to define a server that will be used to verify the revocation status of certificates issues by the PAN-OS device.<br>If you set this privilege to read-only, the administrator can view the **OCSP Responder** configuration for the device but is not allowed to create or edit an OCSP responder configuration. | Yes | Yes | Yes |
| SSL/TLS Service Profile | Controls access to the **SSL/TLS Service Profile** node.<br>If you disable this privilege, the administrator will not see the node or configure a profile that specifies a certificate and a protocol version or range of versions for device services that use SSL/TLS.<br>If you set this privilege to read-only, the administrator can view existing SSL/TLS Service profiles but cannot create or edit them. | Yes | Yes | Yes |
| Response Pages | Controls access to the **Response Pages** node. If you disable this privilege, the administrator will not see the **Response Page** node or be able to define a custom HTML message that is downloaded and displayed instead of a requested web page or file.<br>If you set this privilege to read-only, the administrator can view the **Response Page** configuration for the device but is not allowed to create or edit a response page configuration. | Yes | Yes | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Log Settings | Sets the default state to enable or disable for all of the Log settings described below. | Yes | No | Yes |
| System | Controls access to the **Log Settings > System** node. If you disable this privilege, the administrator will not see the **Log Settings > System** node or be able to specify the severity levels of the system log entries that are logged remotely with Panorama and sent as SNMP traps, syslog messages, and/or email notifications.<br>If you set this privilege to read-only, the administrator can view the **Log Settings > System** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |
| Config | Controls access to the **Log Settings > Config** node. If you disable this privilege, the administrator will not see the **Log Settings > Config** node or be able to specify the configuration log entries that are logged remotely with Panorama, and sent as syslog messages and/or email notification.<br>If you set this privilege to read-only, the administrator can view the **Log Settings > Config** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |
| HIP Match | Controls access to the **Log Settings > HIP Match** node. If you disable this privilege, the administrator will not see the **Log Settings > HIP Match** node or be able to specify the Host Information Profile (HIP) match log settings that are used to provide information on security rules that apply to GlobalProtect clients<br>If you set this privilege to read-only, the administrator can view the **Log Settings > HIP** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |
| Alarms | Controls access to the **Log Settings > Alarms** node. If you disable this privilege, the administrator will not see the **Log Settings > Alarms** node or be able to configure notifications that are generated when a security rule (or group of rules) has been hit repeatedly in a set period of time.<br>If you set this privilege to read-only, the administrator can view the **Log Settings > Alarms** configuration for the device but is not allowed to create or edit a configuration. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Manage Logs | Controls access to the **Log Settings > Manage Logs** node. If you disable this privilege, the administrator will not see the **Log Settings > Manage Logs** node or be able to clear the indicated logs.<br>If you set this privilege to read-only, the administrator can view the **Log Settings > Manage Logs** information but cannot clear any of the logs. | Yes | Yes | Yes |
| Server Profiles | Sets the default state to enable or disable for all of the Server Profiles settings described below. | Yes | No | Yes |
| SNMP Trap | Controls access to the **Server Profiles > SNMP Trap** node. If you disable this privilege, the administrator will not see the **Server Profiles > SNMP Trap** node or be able to specify one or more SNMP trap destinations to be used for system log entries.<br>If you set this privilege to read-only, the administrator can view the **Server Profiles > SNMP Trap Logs** information but cannot specify SNMP trap destinations. | Yes | Yes | Yes |
| Syslog | Controls access to the **Server Profiles > Syslog** node. If you disable this privilege, the administrator will not see the **Server Profiles > Syslog** node or be able to specify one or more syslog servers.<br>If you set this privilege to read-only, the administrator can view the **Server Profiles > Syslog** information but cannot specify syslog servers. | Yes | Yes | Yes |
| Email | Controls access to the **Server Profiles > Email** node. If you disable this privilege, the administrator will not see the **Server Profiles > Email** node or be able to configure an email profile that can be used to enable email notification for system and configuration log entries<br>If you set this privilege to read-only, the administrator can view the **Server Profiles > Email** information but cannot configure and email profile. | Yes | Yes | Yes |
| Netflow | Controls access to the **Server Profiles > Netflow** node. If you disable this privilege, the administrator will not see the **Server Profiles > Netflow** node or be able to define a NetFlow server profile, which specifies the frequency of the export along with the NetFlow servers that will receive the exported data.<br>If you set this privilege to read-only, the administrator can view the **Server Profiles > Netflow** information but cannot define a Netflow profile. | Yes | Yes | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| RADIUS | Controls access to the **Server Profiles > RADIUS** node. If you disable this privilege, the administrator will not see the **Server Profiles > RADIUS** node or be able to configure settings for the RADIUS servers that are identified in authentication profiles.<br>If you set this privilege to read-only, the administrator can view the **Server Profiles > RADIUS** information but cannot configure settings for the RADIUS servers. | Yes | Yes | Yes |
| TACACS+ | Controls access to the **Server Profiles > TACACS+** node.<br>If you disable this privilege, the administrator will not see the node or configure settings for the TACACS+ servers that authentication profiles reference.<br>If you set this privilege to read-only, the administrator can view existing TACACS+ server profiles but cannot add or edit them. | Yes | Yes | Yes |
| LDAP | Controls access to the **Server Profiles > LDAP** node. If you disable this privilege, the administrator will not see the **Server Profiles > LDAP** node or be able to configure settings for the LDAP servers to use for authentication by way of authentication profiles.<br>If you set this privilege to read-only, the administrator can view the **Server Profiles > LDAP** information but cannot configure settings for the LDAP servers. | Yes | Yes | Yes |
| Kerberos | Controls access to the **Server Profiles > Kerberos** node. If you disable this privilege, the administrator will not see the **Server Profiles > Kerberos** node or configure a Kerberos server that allows users to authenticate natively to a domain controller.<br>If you set this privilege to read-only, the administrator can view the **Server Profiles > Kerberos** information but cannot configure settings for Kerberos servers. | Yes | Yes | Yes |
| Local User Database | Sets the default state to enable or disable for all of the Local User Database settings described below. | Yes | No | Yes |
| Users | Controls access to the **Local User Database > Users** node. If you disable this privilege, the administrator will not see the **Local User Database > Users** node or set up a local database on the firewall to store authentication information for remote access users, device administrators, and captive portal users.<br>If you set this privilege to read-only, the administrator can view the **Local User Database > Users** information but cannot set up a local database on the firewall to store authentication information. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| User Groups | Controls access to the **Local User Database > Users** node. If you disable this privilege, the administrator will not see the **Local User Database > Users** node or be able to add user group information to the local database.<br>If you set this privilege to read-only, the administrator can view the **Local User Database > Users** information but cannot add user group information to the local database. | Yes | Yes | Yes |
| Authentication Profile | Controls access to the **Authentication Profile** node. If you disable this privilege, the administrator will not see the **Authentication Profile** node or be able to create or edit authentication profiles that specify local database, RADIUS, TACACS+, LDAP, or Kerberos settings that can be assigned to administrator accounts.<br>If you set this privilege to read-only, the administrator can view the **Authentication Profile** information but cannot create or edit an authentication profile. | Yes | Yes | Yes |
| Authentication Sequence | Controls access to the **Authentication Sequence** node. If you disable this privilege, the administrator will not see the **Authentication Sequence** node or be able to create or edit an authentication sequence.<br>If you set this privilege to read-only, the administrator can view the **Authentication Profile** information but cannot create or edit an authentication sequence. | Yes | Yes | Yes |
| Access Domain | Controls access to the **Access Domain** node. If you disable this privilege, the administrator will not see the **Access Domain** node or be able to create or edit an access domain.<br>If you set this privilege to read-only, the administrator can view the **Access Domain** information but cannot create or edit an access domain. | Yes | Yes | Yes |
| Scheduled Log Export | Controls access to the **Scheduled Log Export** node. If you disable this privilege, the administrator will not see the **Scheduled Log Export** node or be able schedule exports of logs and save them to a File Transfer Protocol (FTP) server in CSV format or use Secure Copy (SCP) to securely transfer data between the device and a remote host.<br>If you set this privilege to read-only, the administrator can view the **Scheduled Log Export Profile** information but cannot schedule the export of logs. | Yes | No | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Software | Controls access to the **Software** node. If you disable this privilege, the administrator will not see the **Software** node or view the latest versions of the PAN-OS software available from Palo Alto Networks, read the release notes for each version, and select a release to download and install.<br>If you set this privilege to read-only, the administrator can view the **Software** information but cannot download or install software. | Yes | Yes | Yes |
| GlobalProtect Client | Controls access to the **GlobalProtect Client** node. If you disable this privilege, the administrator will not see the **GlobalProtect Client** node or view available GlobalProtect releases, download the code or activate the GlobalProtect agent.<br>If you set this privilege to read-only, the administrator can view the available **GlobalProtect Client** releases but cannot download or install the agent software. | Yes | Yes | Yes |
| Dynamic Updates | Controls access to the **Dynamic Updates** node. If you disable this privilege, the administrator will not see the **Dynamic Updates** node or be able to view the latest updates, read the release notes for each update, or select an update to upload and install.<br>If you set this privilege to read-only, the administrator can view the available **Dynamic Updates** releases, read the release notes but cannot upload or install the software. | Yes | Yes | Yes |
| Licenses | Controls access to the **Licenses** node. If you disable this privilege, the administrator will not see the **Licenses** node or be able to view the licenses installed or activate licenses.<br>If you set this privilege to read-only, the administrator can view the installed **Licenses**, but cannot perform license management functions. | Yes | Yes | Yes |
| Support | Controls access to the **Support** node. If you disable this privilege, the administrator will not see the **Support** node or be able to access product and security alerts from Palo Alto Networks or generate tech support or stats dump files.<br>If you set this privilege to read-only, the administrator can view the **Support** node and access product and security alerts but cannot generate tech support or stats dump files. | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Master Key and Diagnostics | Controls access to the **Master Key and Diagnostics** node. If you disable this privilege, the administrator will not see the **Master Key and Diagnostics** node or be able to specify a master key to encrypt private keys on the firewall.<br>If you set this privilege to read-only, the administrator can view the **Master Key and Diagnostics** node and view information about master keys that have been specified but cannot add or edit a new master key configuration. | Yes | Yes | Yes |

## Define User Privacy Settings in the Admin Role Profile

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Privacy | Sets the default state to enable or disable for all of the privacy settings described below. | Yes | N/A | Yes |
| Show Full IP addresses | When set to disable, full IP addresses obtained by traffic running through the Palo Alto firewall are not shown in logs or reports. In place of the IP addresses that are normally displayed, the relevant subnet is displayed.<br> Scheduled reports that are displayed in the interface through **Monitor > Reports** and reports that are sent via scheduled emails will still display full IP addresses. Because of this exception, we recommend that the following settings within the **Monitor** tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler. | Yes | N/A | Yes |
| Show User Names in Logs and Reports | When set to disable, user names obtained by traffic running through the Palo Alto Networks firewall are not shown in logs or reports. Columns where the user names would normally be displayed are empty.<br> Scheduled reports that are displayed in the interface through **Monitor > Reports** or reports that are sent via the email scheduler will still display user names. Because of this exception, we recommend that the following settings within the Monitor tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler. | Yes | N/A | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| View Pcap Files | When set to disable, packet capture files that are normally available within the Traffic, Threat and Data Filtering logs are not displayed. | Yes | N/A | Yes |

## Restrict Admin Access to Commit Functions

**Restrict User Access Using the Commit Setting**

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Commit | When set to disable, an admin cannot commit any changes to a configuration. | Yes | N/A | Yes |

## Restrict Admin Access to Validate Functions

**Restrict User Access Using the Validate Setting**

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Validate | When set to disable, an admin cannot validate a configuration. | Yes | N/A | Yes |

## Provide Granular Access to Global Settings

**Restrict User Access Using the Global Settings**

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Global | Sets the default state to enable or disable for all of the global settings described below. In effect, this setting is only for System Alarms at this time. | Yes | N/A | Yes |
| System Alarms | When set to disable, an admin cannot view or acknowledge alarms that are generated. | Yes | N/A | Yes |

## Provide Granular Access to the Panorama Tab

The following table lists the **Panorama** tab access levels and the custom Panorama administrator roles for which they are available. Firewall administrators cannot access any of these privileges.

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Setup | Specifies whether the administrator can view or edit Panorama setup information, such as **Management**, **Operations**, **Services**, **WildFire**, or **HSM**.<br>If you set the privilege to:<br>• read-only, the administrator can see the information but cannot edit it.<br>• disable this privilege, the administrator cannot see or edit the information. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| High Availability | Specifies whether the administrator can view and manage high availability (HA) settings for the Panorama management server.<br>If you set this privilege to read-only, the administrator can view HA configuration information for the Panorama management server but can't manage the configuration.<br>If you disable this privilege, the administrator can't see or manage HA configuration settings for the Panorama management server. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Config Audit | Specifies whether the administrator can run Panorama configuration audits. If you disable this privilege, the administrator can't run Panorama configuration audits. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Administrators | Specifies whether the administrator can view Panorama administrator account details.<br>You can't enable full access to this function: just read-only access. (Only Panorama administrators with a dynamic role can add, edit, or delete Panorama administrators.) With read-only access, the administrator can see information about his or her own account but no other Panorama administrator accounts.<br>If you disable this privilege, the administrator can't see information about any Panorama administrator account, including his or her own. | Panorama: Yes<br>Device Group/Template: No | No | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Admin Roles | Specifies whether the administrator can view Panorama administrator roles.<br>You can't enable full access to this function: just read-only access. (Only Panorama administrators with a dynamic role can add, edit, or delete custom Panorama roles.) With read-only access, the administrator can see Panorama administrator role configurations but can't manage them.<br>If you disable this privilege, the administrator can't see or manage Panorama administrator roles. | Panorama: Yes<br>Device Group/Template: No | No | Yes | Yes |
| Access Domain | Specifies whether the administrator can view, add, edit, delete, or clone access domain configurations for Panorama administrators. (This privilege controls access only to the configuration of access domains, not access to the device groups, templates, and firewall contexts that are assigned to access domains.)<br>If you set this privilege to read-only, the administrator can view Panorama access domain configurations but can't manage them.<br>If you disable this privilege, the administrator can't see or manage Panorama access domain configurations. | Panorama: Yes<br>Device Group/Template: No<br> You assign access domains to Device Group and Template administrators so they can access the configuration and monitoring data within the device groups, templates, and firewall contexts that are assigned to those access domains. | Yes | Yes | Yes |
| Authentication Profile | Specifies whether the administrator can view, add, edit, delete, or clone authentication profiles for Panorama administrators.<br>If you set this privilege to read-only, the administrator can view Panorama authentication profiles but can't manage them.<br>If you disable this privilege, the administrator can't see or manage Panorama authentication profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Authentication Sequence | Specifies whether the administrator can view, add, edit, delete, or clone authentication sequences for Panorama administrators.<br>If you set this privilege to read-only, the administrator can view Panorama authentication sequences but can't manage them.<br>If you disable this privilege, the administrator can't see or manage Panorama authentication sequences. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Managed Devices | Specifies whether the administrator can view, add, edit, tag, or delete firewalls as managed devices, and install software or content updates on them.<br>If you set this privilege to read-only, the administrator can see managed firewalls but can't add, delete, tag, or install updates on them.<br>If you disable this privilege, the administrator can't view, add, edit, tag, delete, or install updates on managed firewalls.<br>This privilege applies only to the **Panorama > Managed Devices** page. An administrator with Device Deployment privileges can still use the **Panorama > Device Deployment** pages to install updates on managed firewalls. | Panorama: Yes<br>Device Group/Template: Yes | Yes<br>(No for Device Group and Template roles) | Yes | Yes |
| Templates | Specifies whether the administrator can view, edit, add, or delete templates and template stacks.<br>If you set the privilege to read-only, the administrator can see template and stack configurations but can't manage them.<br>If you disable this privilege, the administrator can't see or manage template and stack configurations. | Panorama: Yes<br>Device Group/Template: Yes<br>Device Group and Template administrators can see only the templates and stacks that are within the access domains assigned to those administrators. | Yes<br>(No for Device Group and Template admins) | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Device Groups | Specifies whether the administrator can view, edit, add, or delete device groups.<br>If you set this privilege to read-only, the administrator can see device group configurations but can't manage them.<br>If you disable this privilege, the administrator can't see or manage device group configurations. | Panorama: Yes<br>Device Group/Template: Yes<br>Device Group and Template administrators can access only the device groups that are within the access domains assigned to those administrators. | Yes | Yes | Yes |
| Managed Collectors | Specifies whether the administrator can view, edit, add, or delete managed collectors.<br>If you set this privilege to read-only, the administrator can see managed collector configurations but can't manage them.<br>If you disable this privilege, the administrator can't view, edit, add, or delete managed collector configurations.<br>This privilege applies only to the **Panorama > Managed Collectors** page. An administrator with Device Deployment privileges can still use the **Panorama > Device Deployment** pages to install updates on managed collectors. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Collector Groups | Specifies whether the administrator can view, edit, add, or delete Collector Groups.<br>If you set this privilege to read-only, the administrator can see Collector Groups but can't manage them.<br>If you disable this privilege, the administrator can't see or manage Collector Groups. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| VMware Service Manager | Specifies whether the administrator can view and edit VMware Service Manager settings.<br>If you set this privilege to read-only, the administrator can see the settings but can't perform any related configuration or operational procedures.<br>If you disable this privilege, the administrator can't see the settings or perform any related configuration or operational procedures. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Certificate Management | Sets the default state, enabled or disabled, for all of the Panorama certificate management privileges. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| Certificates | Specifies whether the administrator can view, edit, generate, delete, revoke, renew, or export certificates. This privilege also specifies whether the administrator can import or export HA keys.<br>If you set this privilege to read-only, the administrator can see Panorama certificates but can't manage the certificates or HA keys.<br>If you disable this privilege, the administrator can't see or manage Panorama certificates or HA keys. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Certificate Profile | Specifies whether the administrator can view, add, edit, delete or clone Panorama certificate profiles.<br>If you set this privilege to read-only, the administrator can see Panorama certificate profiles but can't manage them.<br>If you disable this privilege, the administrator can't see or manage Panorama certificate profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| SSL/TLS Service Profile | Specifies whether the administrator can view, add, edit, delete or clone SSL/TLS Service profiles.<br>If you set this privilege to read-only, the administrator can see SSL/TLS Service profiles but can't manage them.<br>If you disable this privilege, the administrator can't see or manage SSL/TLS Service profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Log Settings | Sets the default state, enabled or disabled, for all the log setting privileges. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| System | Specifies whether the administrator can see and configure the settings that control the forwarding of System logs to external services (syslog, email, or SNMP trap servers).<br>If you set this privilege to read-only, the administrator can see the System log forwarding settings but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the settings.<br>On a Panorama M-Series appliance, this privilege pertains only to System logs that Panorama generates. On a Panorama virtual appliance, this privilege applies to System logs that Panorama generates and to System logs that Panorama collects from firewalls. The **Panorama > Collector Groups** page controls the forwarding of System logs that an M-Series appliance collects from firewalls. The **Device > Log Settings** page controls the forwarding of System logs directly from firewalls to external services (without aggregation on Panorama). | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Config | Specifies whether the administrator can see and configure the settings that control the forwarding of Config logs to external services (syslog, email, or SNMP trap servers).<br>If you set this privilege to read-only, the administrator can see the Config log forwarding settings but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the settings.<br> On a Panorama M-Series appliance, this privilege pertains only to Config logs that Panorama generates. On a Panorama virtual appliance, this privilege applies to Config logs that Panorama generates and to Config logs that Panorama collects from firewalls. The **Panorama > Collector Groups** page controls the forwarding of Config logs that an M-Series appliance collects from firewalls. The **Device > Log Settings** page controls the forwarding of Config logs directly from firewalls to external services (without aggregation on Panorama). | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| HIP Match | Specifies whether the administrator can see and configure the settings that control the forwarding of HIP Match logs from a Panorama virtual appliance to external services (syslog, email, or SNMP trap servers).<br>If you set this privilege to read-only, the administrator can see the forwarding settings of HIP Match logs but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the settings.<br> The **Panorama > Collector Groups** page controls the forwarding of HIP Match logs from a Panorama M-Series appliance. The **Device > Log Settings** page controls the forwarding of HIP Match logs directly from firewalls to external services (without aggregation on Panorama). | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Correlation | Specifies whether the administrator can see and configure the settings that control the forwarding of Correlation logs to external services (syslog, email, or SNMP trap servers).<br>If you set this privilege to read-only, the administrator can see the Correlation log forwarding settings but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the settings.<br> The **Panorama > Collector Groups** page controls the forwarding of Correlation logs from a Panorama M-Series appliance. The **Device > Log Settings** page controls the forwarding of Correlation logs directly from firewalls to external services (without aggregation on Panorama). | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Traffic | Specifies whether the administrator can see and configure the settings that control the forwarding of Traffic logs from a Panorama virtual appliance to external services (syslog, email, or SNMP trap servers).<br>If you set this privilege to read-only, the administrator can see the forwarding settings of Traffic logs but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the settings.<br> The **Panorama > Collector Groups** page controls the forwarding of Traffic logs from a Panorama M-Series appliance. The **Objects > Log Forwarding** page controls the forwarding of Traffic logs directly from firewalls to external services (without aggregation on Panorama). | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Threat | Specifies whether the administrator can see and configure the settings that control the forwarding of Threat logs from a Panorama virtual appliance to external services (syslog, email, or SNMP trap servers).<br>If you set this privilege to read-only, the administrator can see the forwarding settings of Threat logs but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the settings.<br> The **Panorama > Collector Groups** page controls the forwarding of Threat logs from a Panorama M-Series appliance. The **Objects > Log Forwarding** page controls the forwarding of Threat logs directly from firewalls to external services (without aggregation on Panorama). | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Wildfire | Specifies whether the administrator can see and configure the settings that control the forwarding of WildFire logs from a Panorama virtual appliance to external services (syslog, email, or SNMP trap servers).<br>If you set this privilege to read-only, the administrator can see the forwarding settings of WildFire logs but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the settings.<br> The **Panorama > Collector Groups** page controls the forwarding of WildFire logs from a Panorama M-Series appliance. The **Objects > Log Forwarding** page controls the forwarding of WildFire logs directly from firewalls to external services (without aggregation on Panorama). | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Server Profiles | Sets the default state, enabled or disabled, for all the server profile privileges.<br> These privileges pertain only to the server profiles that are used for forwarding logs that Panorama generates or collects from firewalls and the server profiles that are used for authenticating Panorama administrators. The **Device > Server Profiles** pages control the server profiles that are used for forwarding logs directly from firewalls to external services (without aggregation on Panorama) and for authenticating firewall administrators. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |
| SNMP Trap | Specifies whether the administrator can see and configure SNMP trap server profiles.<br>If you set this privilege to read-only, the administrator can see SNMP trap server profiles but can't manage them.<br>If you disable this privilege, the administrator can't see or manage SNMP trap server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Syslog | Specifies whether the administrator can see and configure Syslog server profiles.<br>If you set this privilege to read-only, the administrator can see Syslog server profiles but can't manage them.<br>If you disable this privilege, the administrator can't see or manage Syslog server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Email | Specifies whether the administrator can see and configure email server profiles.<br>If you set this privilege to read-only, the administrator can see email server profiles but can't manage them.<br>If you disable this privilege, the administrator can't see or manage email server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| RADIUS | Specifies whether the administrator can see and configure the RADIUS server profiles that are used to authenticate Panorama administrators.<br>If you set this privilege to read-only, the administrator can see the RADIUS server profiles but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the RADIUS server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| TACACS+ | Specifies whether the administrator can see and configure the TACACS+ server profiles that are used to authenticate Panorama administrators.<br>If you disable this privilege, the administrator can't see the node or configure settings for the TACACS+ servers that authentication profiles reference.<br>If you set this privilege to read-only, the administrator can view existing TACACS+ server profiles but can't add or edit them. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| LDAP | Specifies whether the administrator can see and configure the LDAP server profiles that are used to authenticate Panorama administrators.<br>If you set this privilege to read-only, the administrator can see the LDAP server profiles but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the LDAP server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Kerberos | Specifies whether the administrator can see and configure the Kerberos server profiles that are used to authenticate Panorama administrators.<br>If you set this privilege to read-only, the administrator can see the Kerberos server profiles but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the Kerberos server profiles. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Scheduled Config Export | Specifies whether the administrator can view, add, edit, delete, or clone scheduled Panorama configuration exports.<br>If you set this privilege to read-only, the administrator can view the scheduled exports but can't manage them.<br>If you disable this privilege, the administrator can't see or manage the scheduled exports. | Panorama: Yes<br>Device Group/Template: No | Yes | No | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Software | Specifies whether the administrator can: view information about Panorama software updates; download, upload, or install the updates; and view the associated release notes.<br>If you set this privilege to read-only, the administrator can view information about Panorama software updates and view the associated release notes but can't perform any related operations.<br>If you disable this privilege, the administrator can't see Panorama software updates, see the associated release notes, or perform any related operations.<br>This privilege pertains only to software installed on a Panorama management server. The **Panorama > Device Deployment > Software** page controls access to PAN-OS software deployed on firewalls and Panorama software deployed on Dedicated Log Collectors. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Dynamic Updates | Specifies whether the administrator can: view information about Panorama content updates (for example, WildFire updates); download, upload, install, or revert the updates; and view the associated release notes.<br>If you set this privilege to read-only, the administrator can view information about Panorama content updates and view the associated release notes but can't perform any related operations.<br>If you disable this privilege, the administrator can't see Panorama content updates, see the associated release notes, or perform any related operations.<br>This privilege pertains only to content updates installed on a Panorama management server. The **Panorama > Device Deployment > Dynamic Updates** page controls access to content updates deployed on firewalls and Dedicated Log Collectors. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Support | Specifies whether the administrator can: view Panorama support license information, product alerts, and security alerts; activate a support license, generate Tech Support files, and manage cases<br>If you set this privilege to read-only, the administrator can view Panorama support information, product alerts, and security alerts, but can't activate a support license, generate Tech Support files, or manage cases.<br>If you disable this privilege, the administrator can't: see Panorama support information, product alerts, or security alerts; activate a support license, generate Tech Support files, or manage cases. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |
| Device Deployment | Sets the default state, enabled or disabled, for all the device deployment privileges.<br>These privilege pertain only to software and content updates that Panorama administrators deploy on firewalls and Dedicated Log Collectors. The **Panorama > Software** and **Panorama > Dynamic Updates** pages control the software and content updates installed on a Panorama management server. | Panorama: Yes<br>Device Group/Template: Yes | Yes | No | Yes |
| Software | Specifies whether the administrator can: view information about the software updates installed on firewalls and Log Collectors; download, upload, or install the updates; and view the associated release notes.<br>If you set this privilege to read-only, the administrator can see information about the software updates and view the associated release notes but can't deploy the updates to firewalls or dedicated Log Collectors.<br>If you disable this privilege, the administrator can't see information about the software updates, see the associated release notes, or deploy the updates to firewalls or Dedicated Log Collectors. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |

 © Palo Alto Networks, Inc.

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| SSL VPN Client | Specifies whether the administrator can: view information about SSL VPN client software updates on firewalls; download, upload, or activate the updates; and view the associated release notes.<br>If you set this privilege to read-only, the administrator can see information about SSL VPN client software updates and view the associated release notes but can't activate the updates on firewalls.<br>If you disable this privilege, the administrator can't see information about SSL VPN client software updates, see the associated release notes, or activate the updates on firewalls. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |
| GlobalProtect Client | Specifies whether the administrator can: view information about GlobalProtect agent/app software updates on firewalls; download, upload, or activate the updates; and view the associated release notes.<br>If you set this privilege to read-only, the administrator can see information about GlobalProtect agent/app software updates and view the associated release notes but can't activate the updates on firewalls.<br>If you disable this privilege, the administrator can't see information about GlobalProtect agent/app software updates, see the associated release notes, or activate the updates on firewalls. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |

| Access Level | Description | Admin Role Availability | Enable | Read Only | Disable |
|---|---|---|---|---|---|
| Dynamic Updates | Specifies whether the administrator can: view information about the content updates (for example, Applications updates) installed on firewalls and Dedicated Log Collectors; download, upload, or install the updates; and view the associated release notes.<br>If you set this privilege to read-only, the administrator can see information about the content updates and view the associated release notes but can't deploy the updates to firewalls or Dedicated Log Collectors.<br>If you disable this privilege, the administrator can't see information about the content updates, see the associated release notes, or deploy the updates to firewalls or Dedicated Log Collectors. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |
| Licenses | Specifies whether the administrator can view, refresh, and activate firewall licenses.<br>If you set this privilege to read-only, the administrator can view firewall licenses but can't refresh or activate those licenses.<br>If you disable this privilege, the administrator can't view, refresh, or activate firewall licenses. | Panorama: Yes<br>Device Group/Template: Yes | Yes | Yes | Yes |
| Master Key and Diagnostics | Specifies whether the administrator can view and configure a master key by which to encrypt private keys on Panorama.<br>If you set this privilege to read-only, the administrator can view the Panorama master key configuration but can't change it.<br>If you disable this privilege, the administrator can't see or edit the Panorama master key configuration. | Panorama: Yes<br>Device Group/Template: No | Yes | Yes | Yes |

## Panorama Web Interface Access

The custom Panorama administrator roles allow you to define access to the options on Panorama and the ability to only allow access to Device Groups and Templates (**Policies**, **Objects**, **Network**, **Device** tabs).

 © Palo Alto Networks, Inc.

The admin roles you can create are **Panorama** and **Device Group and Template**. You can't assign CLI access privileges to a **Device Group and Template** admin role. If you assign superuser privileges for the CLI to a **Panorama** admin role, administrators with that role can access all features regardless of the web interface privileges you assign.

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Dashboard | Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets. | Yes | No | Yes |
| ACC | Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy > Show Full Ip Addresses** option and/or the **Show User Names In Logs And Reports** option. | Yes | No | Yes |
| Monitor | Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the admin can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab. | Yes | No | Yes |
| Policies | Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the admin can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab. | Yes | No | Yes |
| Objects | Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the admin can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab. | Yes | No | Yes |

| Access Level | Description | Enable | Read Only | Disable |
|---|---|---|---|---|
| Network | Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the admin can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab. | Yes | No | Yes |
| Device | Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any device-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the admin can see, leave the **Device** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab.<br>You can't enable access to the **Admin Roles** or **Administrators** nodes for a role-based administrator even if you enable full access to the **Device** tab. | Yes | No | Yes |
| Panorama | Controls access to the **Panorama** tab. If you disable this privilege, the administrator will not see the **Panorama** tab and will not have access to any Panorama-wide configuration information, such as Managed Devices, Managed Collectors, or Collector Groups.<br>For more granular control over what objects the admin can see, leave the **Panorama** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Panorama Tab. | Yes | No | Yes |
| Validate | When set to disable, an admin cannot validate a configuration. | Yes | N/A | Yes |

 © Palo Alto Networks, Inc.

# Reference: Port Numbers Used by Palo Alto Networks Devices

The following tables list the ports that Palo Alto Networks devices use to communicate with each other, or with other services on the network.


- Ports Used for Management Functions
- Ports Used for HA
- Ports Used for Panorama
- Ports Used for User-ID


## Ports Used for Management Functions

| Destination Port | Protocol | Description |
|---|---|---|
| 22 | TCP | Used for communication from a client system to the firewall CLI interface. |
| 80 | TCP | The port the firewall listens on for Online Certificate Status Protocol (OCSP) updates when acting as an OCSP responder. |
| 123 | UDP | Port the firewall uses for NTP updates. |
| 443 | TCP | Used for communication from a client system to the firewall web interface. This is also the port the firewall and User-ID agent listens on for VM Information source updates.<br>For monitoring an AWS environment, this is the only port that is used.<br>For monitoring a VMware vCenter/ESXi environment, the listening port defaults to 443, but it is configurable. |
| 162 | UDP | Port the firewall, Panorama, or a Log Collector uses to Forward Traps to an SNMP Manager.<br>This port doesn't need to be open on the Palo Alto Networks device. You must configure the Simple Network Management Protocol (SNMP) manager to listen on this port. For details, refer to the documentation of your SNMP management software. |
| 161 | UDP | Port the firewall listens on for polling requests (GET messages) from the SNMP manager. |
| 514<br>514<br>6514 | TCP<br>UDP<br>SSL | Port that the firewall, Panorama, or a Log Collector uses to send logs to a syslog server if you Configure Syslog Monitoring, and the ports that the PAN-OS integrated User-ID agent or Windows-based User-ID agent listens on for authentication syslog messages if you Configure User-ID to Receive User Mappings from a Syslog Sender. |
| 2055 | UDP | Default port the firewall uses to send NetFlow records to a NetFlow collector if you Configure NetFlow Exports, but this is configurable. |

| Destination Port | Protocol | Description |
|---|---|---|
| 5008 | TCP | Port the GlobalProtect Mobile Security Manager listens on for HIP requests from the GlobalProtect gateways.<br>If you are using a third-party MDM system, you can configure the gateway to use a different port as required by the MDM vendor. |
| 6080<br>6081<br>6082 | TCP<br>TCP<br>TCP | Ports used for Captive Portal: 6080 for NT LAN Manager (NTLM) authentication, 6081 for Captive Portal in transparent mode, and 6082 for Captive Portal in redirect mode. |

## Ports Used for HA

Firewalls configured as High Availability (HA) peers must be able to communicate with each other to maintain state information (HA1 control link) and synchronize data (HA2 data link). In Active/Active HA deployments the peer firewalls must also forward packets to the HA peer that owns the session. The HA3 link is a Layer 2 (MAC-in-MAC) link and it does not support Layer 3 addressing or encryption.

| Destination Port | Protocol | Description |
|---|---|---|
| 28769<br>28260 | TCP<br>TCP | Used for the HA1 control link for clear text communication between the HA peer firewalls. The HA1 link is a Layer 3 link and requires an IP address. |
| 28 | TCP | Used for the HA1 control link for encrypted communication (SSH over TCP) between the HA peer firewalls. |
| 28770 | TCP | Listening port for HA1 backup links. |
| 28771 | TCP | Used for heartbeat backups. Palo Alto Networks recommends enabling heartbeat backup on the MGT interface if you use an in-band port for the HA1 or the HA1 backup links. |
| 99<br>29281 | IP<br>UDP | Used for the HA2 link to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between firewalls in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active device (Active/Passive) or active-primary (Active/Active) to the passive device (Active/Passive) or active-secondary (Active/Active). The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.<br>The HA data link can also be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets. |

 © Palo Alto Networks, Inc.

## Ports Used for Panorama

| Destination Port | Protocol | Description |
|---|---|---|
| 22 | TCP | Used for communication from a client system to the Panorama CLI interface. |
| 443 | TCP | Used for communication from a client system to the Panorama web interface. |
| 3978 | TCP | Used for communication between Panorama and managed devices (firewalls and Log Collectors) as well as for communication among Log Collectors in a Collector Group:<br>• For communication between Panorama and firewalls, this is a bi-directional connection on which the firewalls forward logs to Panorama and Panorama pushes configuration changes to the firewalls. Context switching commands are sent over the same connection.<br>• Log Collectors use this destination port to forward logs to Panorama.<br>• For communication with the default Log Collector on an M-Series appliance in Panorama mode and with Dedicated Log Collectors (M-Series appliances in Log Collector mode). |
| 28769 (5.1 and later)<br>28260 (5.0 and later)<br>49160 (5.0 and earlier) | TCP<br>TCP<br>TCP | Used for the HA connectivity and synchronization between Panorama HA peers using clear text communication. Communication can be initiated by either peer. |
| 28 | TCP | Used for the HA connectivity and synchronization between Panorama HA peers using encrypted communication (SSH over TCP). Communication can be initiated by either peer. |
| 28270 (6.0 and later)<br>49190 (5.1 and earlier) | TCP | Used for communication among Log Collectors in a Collector Group for log distribution. |
| 2049 | TCP | Used by the Panorama virtual appliance to write logs to the NFS datastore. |

## Ports Used for User-ID

User-ID is a feature that enables mapping of user IP addresses to usernames and group memberships, enabling user- or group-based policy and visibility into user activity on your network (for example, to be able to quickly track down a user who may be the victim of a threat). To perform this mapping, the firewall, the User-ID agent (either installed on a Windows-based system or the PAN-OS integrated agent running on the firewall), and/or the Terminal Services agent must be able to connect to directory services on your network to perform Group Mapping and User Mapping. Additionally, if the agents are running on systems external to the firewall, they must be able to connect to the firewall to communicate the IP address to username mappings to the firewall. The following table lists the communication requirements for User-ID along with the port numbers required to establish connections.

| Destination Port | Protocol | Description |
|---|---|---|
| 389 | TCP | Port the firewall uses to connect to an LDAP server (plaintext or Start Transport Layer Security (Start TLS) to Map Users to Groups. |
| 3268 | TCP | Port the firewall uses to connect to an Active Directory global catalog server (plaintext or Start TLS) to Map Users to Groups. |
| 636 | TCP | Port the firewall uses for LDAP over SSL connections with an LDAP server to Map Users to Groups. |
| 3269 | TCP | Port the firewall uses for LDAP over SSL connections with an Active Directory global catalog server to Map Users to Groups. |
| 514<br>514<br>6514 | TCP<br>UDP<br>SSL | Port the PAN-OS integrated User-ID agent or Windows-based User-ID agent listens on for authentication syslog messages if you Configure User-ID to Receive User Mappings from a Syslog Sender. |
| 5007 | TCP | Port the firewall listens on for user mapping information from the User-ID or Terminal Services agent. The agent sends the IP address and username mapping along with a timestamp whenever it learns of a new or updated mapping. In addition, it connects to the firewall at regular intervals to refresh known mappings. |
| 5006 | TCP | Port the User-ID agent listens on for User-ID XML API requests. The source for this communication is typically the system running a script that invokes the API. |
| 88 | UDP/TCP | Port the User-ID agent uses to authenticate to a Kerberos server. The device tries UDP first and falls back to TCP. |
| 1812 | UDP | Port the User-ID agent uses to authenticate to a RADIUS server. |
| 49 | TCP | Port the User-ID agent uses to authenticate to a TACACS+ server. |

 © Palo Alto Networks, Inc.

| Destination Port | Protocol | Description |
|---|---|---|
| 135 | TCP | Port the User-ID agent uses to establish TCP-based WMI connections with the Microsoft Remote Procedure Call (RPC) Endpoint Mapper. The Endpoint Mapper then assigns the agent a randomly assigned port in the 49152-65535 port range. The agent uses this connection to make RPC queries for Exchange Server or AD server security logs, session tables. This is also the port used to access Terminal Services.<br>The User-ID agent also uses this port to connect to client systems to perform Windows Management Instrumentation (WMI) probing. |
| 139 | TCP | Port the User-ID agent uses to establish TCP-based NetBIOS connections to the AD server so that it can send RPC queries for security logs and session information.<br>The User-ID agent also uses this port to connect to client systems for NetBIOS probing (supported on the Windows-based User-ID agent only). |
| 445 | TCP | Port the User-ID agent uses to connect to the Active Directory (AD) using TCP-based SMB connections to the AD server for access to user logon information (print spooler and Net Logon). |

# Reset the Firewall to Factory Default Settings

Resetting the firewall to factory defaults will result in the loss of all configuration settings and logs.

| Reset the Firewall to Factory Default Settings | |
|---|---|
| Step 1 Set up a console connection to the firewall. | 1. Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1).<br>If your computer does not have a 9-pin serial port, use a USB-to-serial port connector.<br>2. Enter your login credentials.<br>3. Enter the following CLI command:<br>`debug system maintenance-mode`<br>The firewall will reboot in the maintenance mode. |
| Step 2 Reset the system to factory default settings. | 1. When the device reboots, press `Enter` to continue to the maintenance mode menu.<br> <br>2. Select `Factory Reset` and press `Enter`.<br>3. Select `Factory Reset` and press `Enter` again.<br>The firewall will reboot without any configuration settings. The default username and password to log in to the firewall is admin/admin.<br>To perform initial configuration on the firewall and to set up network connectivity, see Integrate the Firewall into Your Management Network. |

 © Palo Alto Networks, Inc.

# Authentication

Many of the services that Palo Alto Networks devices provide require authentication, including administrator access to the web interface and end user access to Captive Portal, GlobalProtect portals, and GlobalProtect gateways. The authentication methods that you can configure vary by service, and can include Kerberos single sign-on (SSO), external authentication services, certificates and certificate profiles, local database accounts, RADIUS Vendor-Specific Attributes (VSAs), and NT LAN Manager (NTLM).

The following topics describe authentication methods that are common to most device services, procedures to configure them, how to test authentication profiles, and how to troubleshoot authentication issues:

- Configure Kerberos Single Sign-On
- Configure External Authentication
- Test Authentication Server Connectivity
- Troubleshoot Authentication Issues

# Configure Kerberos Single Sign-On

Palo Alto Networks devices support Kerberos V5 single sign-on (SSO) to authenticate administrators to the web interface and end users to Captive Portal. A network that supports Kerberos SSO prompts a user to log in only for initial access to the network (for example, logging in to Microsoft Windows). After this initial login, the user can access any browser-based service in the network (for example, the firewall web interface) without having to log in again until the SSO session expires. (Your Kerberos administrator sets the duration of SSO sessions.) If you enable both Kerberos SSO and external authentication services (for example, a RADIUS server), the device first tries SSO and, only if that fails, falls back to the external service for authentication.

To support Kerberos SSO, your network requires:

- A Kerberos infrastructure, including a key distribution center (KDC) with an authentication server (AS) and ticket-granting service (TGS).
- A Kerberos account for each Palo Alto Networks device that will authenticate users. An account is required to create a Kerberos keytab, which is a file that contains the principal name and hashed password of the device. The SSO process requires the keytab.

| Configure Kerberos Single Sign-On | | |
|---|---|---|
| Step 1 | Create a Kerberos keytab. | 1. Log in to the KDC and open a command prompt.<br>2. Enter the following command, where <principal_name>, <password>, and <algorithm> are variables. The Kerberos principal name and password are of the device, not the user.<br>**ktpass /princ** <principal_name> **/pass** <password> **/crypto** <algorithm> **/ptype KRB5_NT_PRINCIPAL /out** <file_name>**.keytab**<br>If the device is in FIPS or CC mode, the algorithm must be `aes128-cts-hmac-sha1-96` or `aes256-cts-hmac-sha1-96`. Otherwise, you can also use `des3-cbc-sha1` or `arcfour-hmac`. To use an Advanced Encryption Standard (AES) algorithm, the functional level of the KDC must be Windows Server 2008 or later and you must enable AES encryption for the device account.<br>The algorithm in the keytab must match the algorithm in the service ticket that the TGS issues to clients. Your Kerberos administrator determines which algorithms the service tickets use. |
| Step 2 | Import the keytab into an authentication profile. | Configure an authentication profile.<br>1. Enter the **Kerberos Realm** (usually the DNS domain of the users, except that the realm is uppercase).<br>2. **Import** the **Kerberos Keytab** that you created for the device. |
| Step 3 | Assign the authentication profile to the user account or device service. | • Configure an administrator account.<br>• Configure Captive Portal. |

 © Palo Alto Networks, Inc.

# Configure External Authentication

Palo Alto Networks devices can use external servers for many services that require authentication, including administrator access to the web interface and end user access to Captive Portal, GlobalProtect portals and GlobalProtect gateways. The server protocols that Palo Alto Networks devices support include Lightweight Directory Access Protocol (LDAP), Kerberos, Terminal Access Controller Access-Control System Plus (TACACS+), and Remote Authentication Dial-In User Service (RADIUS). If you enable both external authentication and Kerberos single sign-on (SSO), the device first tries SSO and, only if that fails, falls back to the external server for authentication. To configure external authentication, you create an authentication server profile, assign it to an authentication profile, and then enable authentication for an administrator account or device service by assigning the authentication profile to it.

- Configure Authentication Server Profiles
- Configure an Authentication Profile and Sequence
- Enable External Authentication for Users and Services

## Configure Authentication Server Profiles


- Configure a RADIUS Server Profile
- RADIUS Vendor-Specific Attributes for Palo Alto Networks Devices
- Configure a TACACS+ Server Profile
- Configure an LDAP Server Profile
- Configure a Kerberos Server Profile
- Set CHAP and PAP Authentication for RADIUS and TACACS+ Servers


 © Palo Alto Networks, Inc.

## Configure a RADIUS Server Profile

You can configure Palo Alto Networks devices to use a RADIUS server for authenticating users, managing administrator accounts (if they are not local), and collecting RADIUS Vendor-Specific Attributes (VSAs) from GlobalProtect clients. To use a RADIUS server for managing administrator accounts or collecting GlobalProtect clients VSAs, you must define VSAs on the RADIUS server. For details, see the list of supported RADIUS Vendor-Specific Attributes for Palo Alto Networks Devices.

By default, when authenticating to the RADIUS server, the firewall or Panorama first tries Challenge-Handshake Authentication Protocol (CHAP) and falls back to Password Authentication Protocol (PAP) under certain conditions. Optionally, you can override this automatic protocol selection and configure the firewall or Panorama to always use a specific protocol. For details, see Set CHAP and PAP Authentication for RADIUS and TACACS+ Servers.

| Configure a RADIUS Server Profile | |
|---|---|
| Step 1 | Select **Device > Server Profiles > RADIUS** and click **Add**. |
| Step 2 | Enter a **Profile Name** to identify the server profile. |
| Step 3 | For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available. |
| Step 4 | For the **Timeout**, enter an interval in seconds after which an authentication request times out (range is 1-30, default is 3). |
| Step 5 | Enter the number of automatic **Retries** following a **Timeout** before the request fails (range is 1-5, default is 3). |
| Step 6 | For each RADIUS server, click **Add** and enter a **Name** (to identify the server), server IP address or FQDN (**RADIUS Server** field), **Secret**/**Confirm Secret** (a key to encrypt passwords), and server **Port** for authentication requests (default is 1812). |
| Step 7 | Click **OK** and **Commit**. |

## RADIUS Vendor-Specific Attributes for Palo Alto Networks Devices

Palo Alto Networks devices support the following RADIUS Vendor-Specific Attributes (VSAs). To define VSAs on a RADIUS server, you must specify the vendor code (25461 for Palo Alto Networks devices) and the VSA name and number. Some VSAs also require a value.

| Name | Number | Value |
|---|---|---|
| VSAs for administrator account management and authentication | | |
| PaloAlto-Admin-Role | 1 | A default (dynamic) administrative role name or a custom administrative role name on the firewall. |
| PaloAlto-Admin-Access-Domain | 2 | The name of an access domain for firewall administrators (configured in the **Device > Access Domains** page). Define this VSA if the firewall has multiple virtual systems. |
| PaloAlto-Panorama-Admin-Role | 3 | A default (dynamic) administrative role name or a custom administrative role name on Panorama. |
| PaloAlto-Panorama-Admin-Access-Domain | 4 | The name of an access domain for Device Group and Template administrators (configured in the **Panorama > Access Domains** page). |
| PaloAlto-User-Group | 5 | The name of a user group that an authentication profile references. |
| VSAs forwarded from GlobalProtect clients to the RADIUS server | | |
| PaloAlto-User-Domain | 6 | Don't specify a value when you define these VSAs. |
| PaloAlto-Client-Source-IP | 7 | |
| PaloAlto-Client-OS | 8 | |
| PaloAlto-Client-Hostname | 9 | |
| PaloAlto-GlobalProtect-Client-Version | 10 | |

 © Palo Alto Networks, Inc.

## Configure a TACACS+ Server Profile

Terminal Access Controller Access-Control System Plus (TACACS+) protocol provides better Authentication security than RADIUS because it encrypts usernames and passwords (instead of just passwords), and is also more reliable (it uses TCP instead of UDP).

By default, when authenticating to the TACACS+ server, the firewall or Panorama first tries Challenge-Handshake Authentication Protocol (CHAP) and falls back to Password Authentication Protocol (PAP) under certain conditions. Optionally, you can override this automatic protocol selection and configure the firewall or Panorama to always use a specific protocol. For details, see Set CHAP and PAP Authentication for RADIUS and TACACS+ Servers.

| Configure a TACACS+ Server Profile | |
|---|---|
| Step 1 | Select **Device > Server Profiles > TACACS+** and click **Add**. |
| Step 2 | Enter a **Profile Name** to identify the server profile. |
| Step 3 | For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available. |
| Step 4 | For the **Timeout**, enter an interval in seconds after which an authentication request times out (range is 1-20, default is 3). |
| Step 5 | Select the **Use single connection for all authentication** check box to use the same TCP session for all authentications that use this profile. This option improves performance by avoiding the need to start and end a separate TCP session for each authentication. The check box is cleared by default. |
| Step 6 | For each TACACS+ server, click **Add** and enter a **Name** (to identify the server), server IP address or FQDN (**TACACS+ Server** field), **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords), and server **Port** for authentication requests (default is 49). |
| Step 7 | Click **OK** and **Commit**. |

## Configure an LDAP Server Profile

An LDAP server profile enables you to:

- Authenticate administrators and end users of Palo Alto Networks devices.
- Define security rules based on user or group. The LDAP server profile instructs the firewall how to connect and authenticate to the server and how to search the directory for user and group information. You must also configure User-ID to Map Users to Groups. Then you can select users or groups when defining policy rules.

| Configure an LDAP Server Profile | |
|---|---|
| Step 1 | Select **Device > Server Profiles > LDAP** and click **Add**. |
| Step 2 | Enter a **Profile Name** to identify the server profile. |
| Step 3 | For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available. |
| Step 4 | For each LDAP server (up to four), click **Add** and enter a **Name** (to identify the server), server IP address (**LDAP Server** field), and server **Port** (default 389). |
| Step 5 | Select the server **Type** from the drop-down: **active-directory**, **e-directory**, **sun**, or **other**. |
| Step 6 | If you want the device to use SSL or TLS for a more secure connection with the directory server, select the **Require SSL/TLS secured connection** check box (it is selected by default). The protocol that the device uses depends on the server **Port**:<br>• 389 (default)—TLS (Specifically, the device uses the Start TLS operation, which upgrades the initial plaintext connection to TLS.)<br>• 636—SSL<br>• Any other port—The device first tries to use TLS. If the directory server doesn't support TLS, the device falls back to SSL. |
| Step 7 | To improve security, you can select the **Verify Server Certificate for SSL sessions** check box (it is cleared by default) so that the device verifies the certificate that the directory server presents for SSL/TLS connections. If the verification fails, the connection fails. To enable verification, you must also select the **Require SSL/TLS secured connection** check box. The device verifies the certificate in two respects:<br>• The certificate is trusted and valid. For the device to trust the certificate, its root certificate authority (CA) and any intermediate certificates must be in the certificate store under **Device > Certificate Management > Certificates > Device Certificates**. Import the certificate if necessary: see Import a Certificate and Private Key.<br>• The certificate name must match the host **Name** of the LDAP server. The device first checks the certificate attribute Subject AltName for matching, then tries the attribute Subject DN. If the certificate uses the FQDN of the directory server, you must enter that FQDN in the **LDAP Server** field for the name matching to succeed. |
| Step 8 | Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

## Configure a Kerberos Server Profile

A Kerberos server profile enables users to natively authenticate to an Active Directory domain controller or a Kerberos V5-compliant authentication server. This authentication method is interactive, requiring users to enter usernames and passwords, in contrast with Kerberos single sign-on (SSO), which involves transparent authentication.

To use a Kerberos server for authentication, the server must be accessible over an IPv4 address. IPv6 addresses are not supported.

| Configure a Kerberos Server Profile | |
|---|---|
| Step 1 | Select **Device > Server Profiles > Kerberos** and click **Add**. |
| Step 2 | Enter a **Profile Name** to identify the server profile. |
| Step 3 | For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available. |
| Step 4 | For each Kerberos server, click **Add** and enter a **Name** (to identify the server), server IPv4 address or FQDN (**Kerberos Server** field), and an optional **Port** number for communication with the server (default 88). |
| Step 5 | Click **OK** and **Commit**. |

## Set CHAP and PAP Authentication for RADIUS and TACACS+ Servers

When you configure a Palo Alto Networks device to use RADIUS or TACACS+ server authentication for a particular service (for example, Captive Portal), the device first tries to authenticate to the server using Challenge-Handshake Authentication Protocol (CHAP). The device falls back to Password Authentication Protocol (PAP) if the server rejects the CHAP request. This will happen if, for example, the server doesn't support CHAP or isn't configured for CHAP. CHAP is the preferred protocol because it is more secure than PAP. After the device falls back to PAP for a particular RADIUS or TACACS+ server, the device uses only PAP in subsequent attempts to authenticate to that server. PAN-OS records a fall back to PAP as a medium severity event in the System logs. If you modify any fields in the RADIUS or TACACS+ server profile and then commit the changes, the device reverts to first trying CHAP for that server.

In Release 7.0.6 and later releases, you can force the firewall or Panorama to always use a specific protocol for authenticating to the RADIUS or TACACS+ server by entering the following operational CLI command (the `auto` option reverts to the default automatic selection):

```
set authentication radius-auth-type [ auto | chap | pap ]
```

When configuring a RADIUS or TACACS+ server for CHAP, you must define user accounts with reversibly encrypted passwords. Otherwise, CHAP authentication will fail.

 © Palo Alto Networks, Inc.

## Configure an Authentication Profile and Sequence

An authentication profile specifies the authentication service that validates the credentials of an administrator during login and defines how a Palo Alto Networks device accesses the service. If you create a local administrator account on the device, you can authenticate the administrator to the local database, use an external service (RADIUS, TACACS+, LDAP, or Kerberos server), or use Kerberos single sign-on (SSO).

Some networks have multiple databases for different users and user groups. To authenticate to multiple authentication sources (for example, local database and LDAP), configure an authentication sequence. An authentication sequence is a ranked order of authentication profiles that the device matches an administrator against during login. The device checks against the local database first, and then checks each profile in sequence until one successfully authenticates the administrator. An administrator is denied access only if an authentication failure occurs for all the profiles in the authentication sequence.

| **Configure an Authentication Profile and Sequence** | | |
|---|---|---|
| Step 1 | Create a Kerberos keytab.<br><br>Required if the device will use Kerberos SSO authentication. | Create a Kerberos keytab. A keytab is a file that contains Kerberos account information (principal name and hashed password) for the device. |
| Step 2 | Configure an external server profile.<br><br>Required if the device will use an external service for authentication. | • Configure a RADIUS Server Profile.<br>• Configure a TACACS+ Server Profile.<br>• Configure an LDAP Server Profile.<br>• Configure a Kerberos Server Profile. |

**Configure an Authentication Profile and Sequence (Continued)**

Step 3 Configure an authentication profile.

Define one or both of the following authentication phases:

- Kerberos SSO—The device first tries SSO authentication. If that fails, the device falls back to authentication of the **Type** specified in the profile.
- Local database or external authentication—The device prompts the user to enter login credentials, and uses its local database or an external service to authenticate the user.

1. Select **Device > Authentication Profile** and click **Add**.
2. Enter a **Name** to identify the authentication profile.
3. If the firewall has more than one virtual system (vsys), select a **Location** (a vsys or **Shared**) where the profile is available.
4. In the **Authentication** tab, select the authentication **Type**. If you select **RADIUS**, **TACACS+**, **LDAP**, or **Kerberos**, select the authentication **Server Profile** from the drop-down.

   If the **Type** is **LDAP**, define the **Login Attribute**. For Active Directory, enter **sAMAccountName** as the value.
5. (Optional) Specify **User Domain** and **Username Modifier** values to modify the domain/username string that the user will enter during login. This is useful when the authentication service requires the string in a particular format and you don't want to rely on users to correctly enter the domain. Select from the following options:
   - To send only the unmodified user input, leave the **User Domain** blank (the default) and set the **Username Modifier** to the variable **%USERINPUT%** (the default).
   - To prepend a domain to the user input, enter a **User Domain** and set the **Username Modifier** to **%USERDOMAIN%\%USERINPUT%**.
   - To append a domain to the user input, enter a **User Domain** and set the **Username Modifier** to **%USERINPUT%@%USERDOMAIN%**.
6. If you want to enable Kerberos SSO, enter the **Kerberos Realm** (usually the DNS domain of the users, except that the realm is uppercase) and **Import** the **Kerberos Keytab** that you created for the device.
7. Select the **Advanced** tab and, in the Allow List, click **Add** to select the users and groups that can authenticate with this profile. You can select users/groups from the local database or, if you configured an LDAP server profile, from an LDAP-based directory service such as Active Directory. Selecting **all** allows every user to authenticate. By default, the list is empty, meaning no users can authenticate.

   You can also create and allow custom groups based on LDAP filters: see Map Users to Groups.
8. Enter the number of **Failed Attempts** (0-10) to log in that the device allows before locking out the user. The default value 0 means there is no limit.
9. Enter the **Lockout Time** (0-60), which is the number of minutes for which the device locks out the user after reaching the **Failed Attempts** limit. The default value 0 means the lockout applies until an administrator unlocks the user account.
10. Click **OK** to save the authentication profile.


© Palo Alto Networks, Inc.

| Configure an Authentication Profile and Sequence (Continued) | | |
|---|---|---|
| Step 4 | Configure an authentication sequence.<br><br>Required if you want the device to try multiple authentication profiles to authenticate users. The device evaluates the profiles in top-to-bottom order—applying the Kerberos SSO, authentication service, allow list, and account lockout values for each—until one profile successfully authenticates the user. The device denies access only if all the profiles in the sequence fail to authenticate. | 1. Select **Device > Authentication Sequence** and click **Add**.<br>2. Enter a **Name** to identify the authentication sequence.<br>3. If the firewall has more than one virtual system (vsys), select a **Location** (a vsys or **Shared**) where the sequence is available.<br><br>To expedite the authentication process, the best practice is to select the **Use domain to determine authentication profile** check box: the device will match the domain name that a user enters during login with the **User Domain** or **Kerberos Realm** of an authentication profile in the sequence, and then use that profile to authenticate the user. If the device doesn't find a match, or if you clear the check box, the device tries the profiles in the top-to-bottom sequence.<br><br>4. For each authentication profile to include, click **Add** and select from the drop-down. To change the evaluation order of the profiles, select a profile and click **Move Up** or **Move Down**.<br>5. Click **OK** to save the authentication sequence. |
| Step 5 | Assign the authentication profile or sequence. | Assign the authentication profile or sequence to a user or device service. |

## Enable External Authentication for Users and Services

Palo Alto Networks devices can use external services to authenticate administrators, end users, and other devices.

| Enable External Authentication | | |
|---|---|---|
| Step 1 | Configure an external server profile. | • Configure a RADIUS Server Profile.<br>• Configure a TACACS+ Server Profile.<br>• Configure an LDAP Server Profile.<br>• Configure a Kerberos Server Profile. |
| Step 2 | Assign the server profile to an authentication profile.<br>Optionally, you can assign multiple authentication profiles to an authentication sequence. | Configure an Authentication Profile and Sequence. |
| Step 3 | Assign the authentication profile or sequence to a user or device service. | • Administrators: Configure an Administrative Account.<br>• End users:<br>  • Configure Captive Portal.<br>  • Configure the GlobalProtect portal.<br>  • Configure the GlobalProtect gateway.<br>• Device services:<br>  • Configure Routing Information Protocol (RIP).<br>  • Configure Open Shortest Path First (OSPF).<br>  • Configure Border Gateway Protocol (BGP). |


© Palo Alto Networks, Inc.

# Test Authentication Server Connectivity

After you configure an authentication profile on a Palo Alto Networks firewall or Panorama manager, you can use the test authentication feature to determine if the device can communicate with the back-end authentication server and if the authentication request was successful. You can additionally test authentication profiles used for GlobalProtect and Captive Portal authentication. You can perform authentication tests on the candidate configuration, so that you know the configuration is correct before committing.

Authentication server connectivity testing is supported for local database, RADIUS, TACACS+, LDAP, and Kerberos authentication.

The following topics describe how to use the test authentication command and provides use case examples:

- Run the Test Authentication Command
- Local Database Authentication Profile Use Case
- RADIUS Authentication Profile Use Case
- TACACS+ Authentication Profile Use Case
- LDAP Authentication Profile Use Case
- Kerberos Authentication Profile Use Case

© Palo Alto Networks, Inc.

# Run the Test Authentication Command

| Run the Test Authentication Command | |
|---|---|
| Step 1 | On the PAN-OS firewall or Panorama server, Configure an authentication profile. You do not need to commit the authentication or server profile configuration prior to testing. |
| Step 2 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall. |
| Step 3 | (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.<br>This is required on firewalls with multiple virtual systems (vsys) configured, so the test authentication command can locate the user (Global Protect or Captive Portal, for example) in the correct vsys.<br>To define the target vsys:<br>`admin@PA-3060> set system setting target-vsys <vsys-name>`<br>For example, if the user is defined in vsys2, run the following command:<br>`admin@PA-3060> set system setting target-vsys vsys2`<br>The `target-vsys` command is per-login session, so the system clears the option when you log off. |
| Step 4 | Test an authentication profile by entering the following command:<br>`admin@PA-3060> test authentication authentication-profile <authentication-profile-name> username <username> password`<br>For example, to test an authentication profile named my-profile for a user named bsimpson, run the following command:<br>`admin@PA-3060> test authentication authentication-profile my-profile username bsimpson password`<br>When entering authentication profile names and server profile names in the test command, the names are case sensitive. Also, if the authentication profile has a username modifier defined, you must enter the modifier with the username. For example, if you add the username modifier %USERINPUT%@%USERDOMAIN% for a user named bsimpson and the domain name is mydomain.com, enter bsimpson@mydomain.com as the username. This will ensure that the correct credentials are sent to the authentication server. In this example, mydomain.com is the domain that you define in the User Domain field in the Authentication profile. |
| Step 5 | View the output of the test results.<br>If the authentication profile is configured correctly, the output displays `Authentication succeeded`. If there is a configuration issue, the output displays information to help you troubleshoot the configuration.<br>For example use cases on the supported authentication profile types, see Local Database Authentication Profile Use Case.<br>The output results vary based on several factors related to the authentication type that you are testing as well as the type of issue. For example, RADIUS and TACACS+ use different underlying libraries, so the same issue that exists for both of these types will produce different errors. Also, if there is a network problem, such as using an incorrect port or IP address in the authentication server profile, the output error is not specific. This is because the test command cannot perform the initial handshake between the firewall and the authentication server to determine details about the issue. |

 © Palo Alto Networks, Inc.

## Local Database Authentication Profile Use Case

The following example shows how to test a Local Database authentication profile named LocalDB for a user named User1-LocalDB and how to troubleshoot error conditions that arise. For details on using the test authentication command, see Run the Test Authentication Command.

| Local Database Authentication Profile Test Example | |
|---|---|
| Step 1 | On the PAN-OS firewall, ensure that you have an administrator configured with the type Local Database. For information on administrator accounts, refer to Manage Firewall Administrators. |
| Step 2 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall. |
| Step 3 | (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.<br>This is required on firewalls with multiple virtual systems (vsys) configured, so the test authentication command can locate the user (Global Protect or Captive Portal, for example) in the correct vsys.<br>To define the target vsys:<br>`admin@PA-3060> set system setting target-vsys <vsys-name>`<br>For example, if the user is defined in vsys2, run the following command:<br>`admin@PA-3060> set system setting target-vsys vsys2`<br>The `target-vsys` command is per-login session, so the system clears the option when you log off. |
| Step 4 | Run the following CLI command:<br>`admin@PA-3060> test authentication authentication-profile LocalDB-Profile username User1-LocalDB password` |
| Step 5 | When prompted, enter the password for the User1-LocalDB account. The following output shows that the test failed:<br>`Allow list check error:`<br>`Do allow list check before sending out authentication request...`<br>`User User1-LocalDB is not allowed with authentication profile LocalDB-Profile`<br>In this case, the last line of the output shows that the user is not allowed, which indicates a configuration problem in the authentication profile. |
| Step 6 | To resolve this issue, modify the authentication profile and add the user to the Allow List.<br>1. On the firewall, select **Device > Authentication Profile** and modify the profile named LocalDB-Profile.<br>2. Click the **Advanced** tab and add User1-LocalDB to the Allow List.<br>3. Click **OK** to save the change. |
| Step 7 | Run the test command again. The following output shows that the test is successful:<br>`Do allow list check before sending out authentication request...`<br>`name "User1-LocalDB" has an exact match in allow list`<br>`Authentication by Local User Database for user "User1-LocalDB"`<br>`Authentication succeeded for Local User Database user "User1-LocalDB"` |

## RADIUS Authentication Profile Use Case

The following example shows how to test a RADIUS profile named RADIUS-Profile for a user named User2-RADIUS and how to troubleshoot error conditions that arise. For details on using the test authentication command, see Run the Test Authentication Command.

| RADIUS Authentication Profile Test Example | |
|---|---|
| Step 1 | On the PAN-OS firewall, Configure a RADIUS Server Profile and Configure an authentication profile. In the authentication profile, you select the new RADIUS server profile in the **Server Profile** drop-down. |
| Step 2 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall. |
| Step 3 | (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.<br>This is required on firewalls with multiple virtual systems (vsys) configured, so the test authentication command can locate the user (Global Protect or Captive Portal, for example) in the correct vsys.<br>To define the target vsys:<br>`admin@PA-3060> set system setting target-vsys <vsys-name>`<br>For example, if the user is defined in vsys2, run the following command:<br>`admin@PA-3060> set system setting target-vsys vsys2`<br>The `target-vsys` command is per-login session, so the system clears the option when you log off. |
| Step 4 | Run the following CLI command:<br>`admin@PA-3060> test authentication authentication-profile RADIUS-Profile username User2-RADIUS password` |
| Step 5 | When prompted, enter the password for the User2-RADIUS account. The following output shows that the test failed:<br>`Do allow list check before sending out authentication request...`<br>`name "User2-RADIUS" is in group "all"`<br>`Authentication to RADIUS server at 10.5.104.99:1812 for user "User2-RADIUS"`<br>`Egress: 10.5.104.98`<br>`Authentication type: CHAP`<br>`Now send request to remote server ...`<br>`RADIUS error: Invalid RADIUS response received - Bad MD5`<br>`Authentication failed against RADIUS server at 10.5.104.99:1812 for user "User2-RADIUS"`<br>In this case, the output shows `Bad MD5`, which indicates that there may be an issue with the secret defined in the RADIUS server profile. |

 © Palo Alto Networks, Inc.

**RADIUS Authentication Profile Test Example**

Step 6 To resolve this issue, modify the RADIUS server profile and ensure that the secret defined on the RADIUS server matches the secret in the server profile.

1. On the firewall, select **Device > Server Profiles > RADIUS** and modify the profile named RADIUS-Profile.
2. In the Servers section, locate the RADIUS server and modify the **Secret** field.
3. Type in the correct secret and then retype to confirm.
4. Click **OK** to save the change.

Step 7 Run the test command again. The following output shows that the test is successful:

```
Do allow list check before sending out authentication request...

name "User2-RADIUS" is in group "all"

Authentication to RADIUS server at 10.5.104.99:1812 for user "User2-RADIUS"

Egress: 10.5.104.98

Authentication type: CHAP

Now send request to remote server ...

RADIUS CHAP auth request is NOT accepted, try PAP next

Authentication type: PAP

Now send request to remote server ...

Authentication succeeded against RADIUS server at 10.5.104.99:1812 for user "User2-RADIUS"

Authentication succeeded for user "User2-RADIUS"
```

## TACACS+ Authentication Profile Use Case

The following example shows how to test a TACACS+ profile named TACACS-Profile for a user named User3-TACACS and how to troubleshoot error conditions that arise. For details on using the test authentication command, see Run the Test Authentication Command.

| TACACS+ Authentication Profile Test Example | |
|---|---|
| Step 1 | On the PAN-OS firewall, Configure a TACACS+ Server Profile and Configure an authentication profile. In the authentication profile, you select the new TACACS+ server profile in the **Server Profile** drop-down. |
| Step 2 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall. |
| Step 3 | (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.<br>This is required on firewalls with multiple virtual systems (vsys) configured, so the test authentication command can locate the user (Global Protect or Captive Portal, for example) in the correct vsys.<br>To define the target vsys:<br>`admin@PA-3060> set system setting target-vsys <vsys-name>`<br>For example, if the user is defined in vsys2, run the following command:<br>`admin@PA-3060> set system setting target-vsys vsys2`<br>The `target-vsys` command is per-login session, so the system clears the option when you log off. |
| Step 4 | Run the following CLI command:<br>`admin@PA-3060> test authentication authentication-profile TACACS-Profile username User3-TACACS password` |

 © Palo Alto Networks, Inc.

**TACACS+ Authentication Profile Test Example**

Step 5 When prompted, enter the password for the User3-TACASC account. The following output shows that the test failed:

```
Do allow list check before sending out authentication request...

name "User2-TACACS" is in group "all"

Authentication to TACACS+ server at '10.5.196.62' for user 'User2-TACACS'

Server port: 49, timeout: 30, flag: 0

Egress: 10.5.104.98

Attempting CHAP authentication ...

CHAP authentication request is created

Sending credential: xxxxxx

Failed to send CHAP authentication request: Network read timed out

Attempting PAP authentication ...

PAP authentication request is created

Failed to send PAP authentication request: Network read timed out

Returned status: -1

Authentication failed against TACACS+ server at 10.5.196.62:49 for user User2-TACACS

Authentication failed for user "User2-TACACS"
```

The output shows error `Network read timed out,` which indicates that the TACACS+ server could not decrypt the authentication request. In this case, there may be an issue with the secret defined in the TACACS+ server profile.

Step 6 To resolve this issue, modify the TACACS+ server profile and ensure that the secret defined on the TACACS+ server matches the secret in the server profile.

1. On the firewall, select **Device > Server Profiles > TACACS+** and modify the profile named TACACS-Profile.
2. In the Servers section, locate the TACACS+ server and modify the **Secret** field.
3. Type in the correct secret and then retype to confirm.
4. Click **OK** to save the change.

**TACACS+ Authentication Profile Test Example**

Step 7 Run the test command again. The following output shows that the test is successful:

```
Do allow list check before sending out authentication request...
name "User2-TACACS" is in group "all"
Authentication to TACACS+ server at '10.5.196.62' for user 'User2-TACACS'
Server port: 49, timeout: 30, flag: 0
Egress: 10.5.104.98
Attempting CHAP authentication ...
CHAP authentication request is created
Sending credential: xxxxxx
CHAP authentication request is sent
Authentication succeeded!
Authentication succeeded for user "User2-TACACS"
```

 © Palo Alto Networks, Inc.

## LDAP Authentication Profile Use Case

The following example shows how to test a LDAP authentication profile named LDAP-Profile for a user named User4-LDAP and how to troubleshoot error conditions that arise. For details on using the test authentication command, see Run the Test Authentication Command.

| LDAP Authentication Profile Test Example | |
|---|---|
| Step 1 | On the PAN-OS firewall, Configure an LDAP Server Profile and Configure an authentication profile. In the authentication profile, you select the new LDAP server profile in the **Server Profile** drop-down. |
| Step 2 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall. |
| Step 3 | (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.<br>This is required on firewalls with multiple virtual systems (vsys) configured, so the test authentication command can locate the user (Global Protect or Captive Portal, for example) in the correct vsys.<br>To define the target vsys:<br>`admin@PA-3060> set system setting target-vsys <vsys-name>`<br>For example, if the user is defined in vsys2, run the following command:<br>`admin@PA-3060> set system setting target-vsys vsys2`<br>The `target-vsys` command is per-login session, so the system clears the option when you log off. |
| Step 4 | Run the following CLI command:<br>`admin@PA-3060> test authentication authentication-profile LDAP-Profile username User4-LDAP password` |
| Step 5 | When prompted, enter the password for the User4-LDAP account. The following output shows that the test failed:<br>`Do allow list check before sending out authentication request...`<br>`name "User4-LDAP" is in group "all"`<br>`Authentication to LDAP server at 10.5.104.99 for user "User4-LDAP"`<br>`Egress: 10.5.104.98`<br>`Type of authentication: plaintext`<br>`Starting LDAP connection...`<br>`Succeeded to create a session with LDAP server`<br>`parse error of dn and attributes for user "User4-LDAP"`<br>`Authentication failed against LDAP server at 10.5.104.99:389 for user "User4-LDAP"`<br>`Authentication failed for user "User4-LDAP"`<br>The output shows `parse error of dn and attributes for user User4-LDAP,` which indicates a BIND DN value issues in the LDAP server profile. In this case, a Domain Component (DC) value is incorrect. |

**LDAP Authentication Profile Test Example**

Step 6 To resolve this issue, modify the LDAP server profile and ensure that the Bind DN DC value is correct by comparing the DC value with the DC value of the LDAP server.

1. On the firewall, select **Device > Server Profiles > LDAP** and modify the profile named LDAP-Profile.
2. In the Server settings section, enter the correct value for the DC in the `Bind DN` field. In this case, the correct value for the DC is MGMT-GROUP
3. Click **OK** to save the change.

Step 7 Run the test command again. The following output shows that the test is successful:

```
Do allow list check before sending out authentication request...

name "User4-LDAP" is in group "all"

Authentication to LDAP server at 10.5.104.99 for user "User4-LDAP"

Egress: 10.5.104.98

Type of authentication: plaintext

Starting LDAP connection...

Succeeded to create a session with LDAP server

DN sent to LDAP server: CN=User4-LDAP,CN=Users,DC=MGMT-GROUP,DC=local

User expires in days: never

Authentication succeeded for user "User4-LDAP"
```

 © Palo Alto Networks, Inc.

## Kerberos Authentication Profile Use Case

The following example shows how to test a Kerberos profile named Kerberos-Profile for a user named User5-Kerberos and how to troubleshoot error conditions that arise. For details on using the test authentication command, see Run the Test Authentication Command.

| Kerberos Authentication Profile Test Example | |
|---|---|
| Step 1 | On the PAN-OS firewall, Configure a Kerberos Server Profile and Configure an authentication profile. In the authentication profile, you select the new Kerberos server profile in the **Server Profile** drop-down. |
| Step 2 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall. |
| Step 3 | (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.<br>This is required on firewalls with multiple virtual systems (vsys) configured, so the test authentication command can locate the user (Global Protect or Captive Portal, for example) in the correct vsys.<br>To define the target vsys:<br>`admin@PA-3060> set system setting target-vsys <vsys-name>`<br>For example, if the user is defined in vsys2, run the following command:<br>`admin@PA-3060> set system setting target-vsys vsys2`<br>The `target-vsys` command is per-login session, so the system clears the option when you log off. |
| Step 4 | Run the following CLI command:<br>`admin@PA-3060> test authentication authentication-profile Kerberos-Profile username User5-Kerberos password` |
| Step 5 | When prompted, enter the password for the User5-Kerberos account. The following output shows that the test failed:<br>`Do allow list check before sending out authentication request...`<br>`name "User5-Kerberos" is in group "all"`<br>`Authentication to KERBEROS server at '10.5.104.99' for user 'User5-Kerberos'`<br>`Realm: 'Bad-MGMT-GROUP.LOCAL'`<br>`Egress: 10.5.104.98`<br>`KERBEROS configuration file is created`<br>`KERBEROS authcontext is created. Now authenticating ...`<br>`Kerberos principal is created`<br>`Sending authentication request to KDC...`<br>`Authentication failure: Wrong realm: 'Bad-MGMT-GROUP.LOCAL' (code: -1765328316)`<br>`Authentication failed against KERBEROS server at 10.5.104.99:88 for user "User5-Kerberos"`<br>`Authentication failed for user "User5-Kerberos"`<br>In this case, the output shows `Wrong realm`, which indicates that the Kerberos realm has an incorrect value. |

**Kerberos Authentication Profile Test Example**

Step 6 To resolve this issue, modify the Kerberos server profile and ensure that the Realm value is correct by comparing the realm name on the Kerberos server.

1. On the firewall, select **Device > Authentication Profiles** and modify the profile named Kerberos-Profile.
2. In the Kerberos Realm field, enter the correct value. In this case, the correct realm is mgmt-group.local.
3. Click **OK** to save the change.

Step 7 Run the test command again. The following output shows that the test is successful:

```
Do allow list check before sending out authentication request...

name "User5-Kerberos" is in group "all"

Authentication to KERBEROS server at '10.5.104.99' for user 'User5-Kerberos'

Realm: 'MGMT-GROUP.LOCAL'

Egress: 10.5.104.98

KERBEROS configuration file is created

KERBEROS authcontext is created. Now authenticating ...

Kerberos principal is created

Sending authentication request to KDC...

Authentication succeeded!

Authentication succeeded for user "User5-Kerberos"
```

 © Palo Alto Networks, Inc.

# Troubleshoot Authentication Issues

When users fail to authenticate to a Palo Alto Networks device, or the Authentication process takes longer than expected, analyzing authentication-related information can help you determine whether the failure or delay resulted from:

- User behavior—For example, users are locked out after entering the wrong credentials or a high volume of users are simultaneously attempting access.
- System or network issues—For example, an authentication server is inaccessible.
- Configuration issues—For example, the Allow List of an authentication profile doesn't have all the users it should have.

The following CLI commands display information that can help you troubleshoot these issues:

| Task | Command |
|---|---|
| Use the `show authentication locked-users` command to display the number of locked user accounts associated with the authentication profile (`auth-profile` option), authentication sequence (`is-seq` option), or virtual system (`vsys` option).<br> To unlock administrator users, use the command: **`set management-server unlock admin`** `<user_name>` | `show authentication locked-users`<br>`{`<br>`vsys <value> \|`<br>`auth-profile <value> \|`<br>`is-seq`<br>`{yes \| no}`<br>`{auth-profile \| vsys} <value>`<br>`}` |

| Task | Command |
|---|---|
| Use the `debug authentication` command to troubleshoot authentication events.<br>Use the `show` options to display authentication request statistics and the current debugging level:<br>• `show` displays the current debugging level for the authentication service (authd).<br>• `show-active-requests` displays the number of active checks for authentication requests, allow lists, and locked user accounts.<br>• `show-pending-requests` displays the number of pending checks for authentication requests, allow lists, and locked user accounts.<br>• `connection-show` displays authentication request and response statistics for all authentication servers or for a specific protocol type.<br>Use the `connection-debug` options to enable or disable authentication debugging:<br>• Use the `on` option to enable or the `off` option to disable debugging for authd.<br>• Use the `connection-debug-on` option to enable or the `connection-debug-off` option to disable debugging for all authentication servers or for a specific protocol type. | `debug authentication`<br>`{`<br>`on {debug \| dump \| error \| info \| warn} \|`<br>`show \|`<br>`show-active-requests \|`<br>`show-pending-requests \|`<br>`connection-show \|`<br>`{`<br>`connection-id \|`<br>`protocol-type`<br>`{`<br>`Kerberos connection-id <value> \|`<br>`LDAP connection-id <value> \|`<br>`RADIUS connection-id <value> \|`<br>`TACACS+ connection-id <value> \|`<br>`}`<br>`connection-debug-on \|`<br>`{`<br>`connection-id \|`<br>`debug-prefix \|`<br>`protocol-type`<br>`{`<br>`Kerberos connection-id <value> \|`<br>`LDAP connection-id <value> \|`<br>`RADIUS connection-id <value> \|`<br>`TACACS+ connection-id <value> \|`<br>`}`<br>`connection-debug-off \|`<br>`{`<br>`connection-id \|`<br>`protocol-type`<br>`{`<br>`Kerberos connection-id <value> \|`<br>`LDAP connection-id <value> \|`<br>`RADIUS connection-id <value> \|`<br>`TACACS+ connection-id <value> \|`<br>`}`<br>`connection-debug-on`<br>`}` |

 © Palo Alto Networks, Inc.

# Certificate Management

The following topics describe the different keys and certificates that Palo Alto Networks® devices use, and how to obtain and manage them:

- Keys and Certificates
- Certificate Revocation
- Certificate Deployment
- Set Up Verification for Certificate Revocation Status
- Configure the Master Key
- Obtain Certificates
- Export a Certificate and Private Key
- Configure a Certificate Profile
- Configure an SSL/TLS Service Profile
- Configure the Key Size for SSL Forward Proxy Server Certificates
- Revoke and Renew Certificates
- Secure Keys with a Hardware Security Module

# Keys and Certificates

To ensure trust between parties in a secure communication session, Palo Alto Networks devices use digital certificates. Each certificate contains a cryptographic key to encrypt plaintext or decrypt cyphertext. Each certificate also includes a digital signature to authenticate the identity of the issuer. The issuer must be in the list of trusted certificate authorities (CAs) of the authenticating party. Optionally, the authenticating party verifies the issuer did not revoke the certificate (see Certificate Revocation).

Palo Alto Networks devices use certificates in the following applications:

- User authentication for Captive Portal, GlobalProtect™, Mobile Security Manager, and web interface access to a Palo Alto Networks device.
- Device authentication for GlobalProtect VPN (remote user-to-site or large scale).
- Device authentication for IPSec site-to-site VPN with Internet Key Exchange (IKE).
- Decrypting inbound and outbound SSL traffic.

  A firewall decrypts the traffic to apply policy rules, then re-encrypts it before forwarding the traffic to the final destination. For outbound traffic, the firewall acts as a forward proxy server, establishing an SSL/TLS connection to the destination server. To secure a connection between itself and the client, the firewall uses a *signing certificate* to automatically generate a copy of the destination server certificate.

The following table describes the keys and certificates that Palo Alto Networks devices use. As a best practice, use different keys and certificates for each usage.

**Table: Palo Alto Networks Device Keys/Certificates**

| Key/Certificate Usage | Description |
|---|---|
| Administrative Access | Secure access to device administration interfaces (HTTPS access to the web interface) requires a server certificate for the MGT interface (or a designated interface on the dataplane if the device does not use MGT) and, optionally, a certificate to authenticate the administrator. |
| Captive Portal | In deployments where Captive Portal identifies users who access HTTPS resources, designate a server certificate for the Captive Portal interface. If you configure Captive Portal to use certificates (instead of, or in addition to, username/password credentials) for user identification, designate a user certificate also. For more information on Captive Portal, see Map IP Addresses to Usernames Using Captive Portal. |
| Forward Trust | For outbound SSL/TLS traffic, if a firewall acting as a forward proxy trusts the CA that signed the certificate of the destination server, the firewall uses the forward trust CA certificate to generate a copy of the destination server certificate to present to the client. To set the key size, see Configure the Key Size for SSL Forward Proxy Server Certificates. For added security, store the key on a hardware security module (for details, see Secure Keys with a Hardware Security Module). |
| Forward Untrust | For outbound SSL/TLS traffic, if a firewall acting as a forward proxy does not trust the CA that signed the certificate of the destination server, the firewall uses the forward untrust CA certificate to generate a copy of the destination server certificate to present to the client. |

 © Palo Alto Networks, Inc.

| Key/Certificate Usage | Description |
|---|---|
| SSL Inbound Inspection | The keys that decrypt inbound SSL/TLS traffic for inspection and policy enforcement. For this application, import onto the firewall a private key for each server that is subject to SSL/TLS inbound inspection. See Configure SSL Inbound Inspection. |
| SSL Exclude Certificate | Certificates for servers to exclude from SSL/TLS decryption. For example, if you enable SSL decryption but your network includes servers for which the firewall should not decrypt traffic (for example, web services for your HR systems), import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Configure Decryption Exceptions. |
| GlobalProtect | All interaction among GlobalProtect components occurs over SSL/TLS connections. Therefore, as part of the GlobalProtect deployment, deploy server certificates for all GlobalProtect portals, gateways, and Mobile Security Managers. Optionally, deploy certificates for authenticating users also.<br>Note that the GlobalProtect Large Scale VPN (LSVPN) feature requires a CA signing certificate. |
| Site-to-Site VPNs (IKE) | In a site-to-site IPSec VPN deployment, peer devices use Internet Key Exchange (IKE) gateways to establish a secure channel. IKE gateways use certificates or preshared keys to authenticate the peers to each other. You configure and assign the certificates or keys when defining an IKE gateway on a firewall. See Site-to-Site VPN Overview. |
| Master Key | The firewall uses a master key to encrypt all private keys and passwords. If your network requires a secure location for storing private keys, you can use an encryption (wrapping) key stored on a hardware security module (HSM) to encrypt the master key. For details, see Encrypt a Master Key Using an HSM. |
| Secure Syslog | The certificate to enable secure connections between the firewall and a syslog server. See Syslog Field Descriptions. |
| Trusted Root CA | The designation for a root certificate issued by a CA that the firewall trusts. The firewall can use a self-signed root CA certificate to automatically issue certificates for other applications (for example, SSL Forward Proxy).<br>Also, if a firewall must establish secure connections with other firewalls, the root CA that issues their certificates must be in the list of trusted root CAs on the firewall. |

# Certificate Revocation

Palo Alto Networks devices use digital certificates to ensure trust between parties in a secure communication session. Configuring a device to check the revocation status of certificates provides additional security. A party that presents a revoked certificate is not trustworthy. When a certificate is part of a chain, the device checks the status of every certificate in the chain except the root CA certificate, for which the device cannot verify revocation status.

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority that issued the certificate must revoke it.

Palo Alto Networks devices support the following methods for verifying certificate revocation status. If you configure both, the devices first try the OCSP method; if the OCSP server is unavailable, the devices use the CRL method.

- Certificate Revocation List (CRL)
- Online Certificate Status Protocol (OCSP)

In PAN-OS, certificate revocation status verification is an optional feature. It is a best practice to enable it for certificate profiles, which define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPSec VPN, and web interface access to Palo Alto Network devices.

## Certificate Revocation List (CRL)

Each certificate authority (CA) periodically issues a certificate revocation list (CRL) to a public repository. The CRL identifies revoked certificates by serial number. After the CA revokes a certificate, the next CRL update will include the serial number of that certificate.

The Palo Alto Networks firewall downloads and caches the last-issued CRL for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the CRL for the issuing CA. Also, the cache only stores a CRL until it expires.

The firewall supports CRLs only in Distinguished Encoding Rules (DER) format. If the firewall downloads a CRL in any other format—for example, Privacy Enhanced Mail (PEM) format—any revocation verification process that uses that CRL will fail when a user performs an activity that triggers the process (for example, sending outbound SSL data). The firewall will generate a system log for the verification failure. If the verification was for an SSL certificate, the firewall will also display the SSL Certificate Errors Notify response page to the user.

To use CRLs for verifying the revocation status of certificates used for the decryption of inbound and outbound SSL/TLS traffic, see Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

To use CRLs for verifying the revocation status of certificates that authenticate users and devices, configure a certificate profile and assign it to the interfaces that are specific to the application: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, or web interface access to Palo Alto Networks devices. For details, see Configure Revocation Status Verification of Certificates Used for User/Device Authentication.

 © Palo Alto Networks, Inc.

## Online Certificate Status Protocol (OCSP)

When establishing an SSL/TLS session, clients can use Online Certificate Status Protocol (OCSP) to check the revocation status of the authentication certificate. The authenticating client sends a request containing the serial number of the certificate to the OCSP responder (server). The responder searches the database of the certificate authority (CA) that issued the certificate and returns a response containing the status (*good*, *revoked* or *unknown*) to the client. The advantage of the OCSP method is that it can verify status in real-time, instead of depending on the issue frequency (hourly, daily, or weekly) of CRLs.

The Palo Alto Networks firewall downloads and caches OCSP status information for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the OCSP information for the issuing CA. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall as an OCSP responder (see Configure an OCSP Responder).

To use OCSP for verifying the revocation status of certificates when the firewall functions as an SSL forward proxy, perform the steps under Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

The following applications use certificates to authenticate users and/or devices: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, and web interface access to Palo Alto Networks devices. To use OCSP for verifying the revocation status of the certificates:

- Configure an OCSP responder.
- Enable the HTTP OCSP service on the firewall.
- Create or obtain a certificate for each application.
- Configure a certificate profile for each application.
- Assign the certificate profile to the relevant application.

To cover situations where the OCSP responder is unavailable, configure CRL as a fall-back method. For details, see Configure Revocation Status Verification of Certificates Used for User/Device Authentication.

# Certificate Deployment

The basic approaches to deploy certificates for Palo Alto Networks devices are:

- **Obtain certificates from a trusted third-party CA**—The benefit of obtaining a certificate from a trusted third-party certificate authority (CA) such as VeriSign or GoDaddy is that end clients will already trust the certificate because common browsers include root CA certificates from well-known CAs in their trusted root certificate stores. Therefore, for applications that require end clients to establish secure connections with a Palo Alto Network device, purchase a certificate from a CA that the end clients trust to avoid having to pre-deploy root CA certificates to the end clients. (Some such applications are a GlobalProtect portal or GlobalProtect Mobile Security Manager.) However, note that most third-party CAs cannot issue signing certificates. Therefore, this type of certificate is not appropriate for applications (for example, SSL/TLS decryption and large-scale VPN) that require the firewall to issue certificates. See Obtain a Certificate from an External CA.

- **Obtain certificates from an enterprise CA**—Enterprises that have their own internal CA can use it to issue certificates for firewall applications and import them onto the firewall. The benefit is that end clients probably already trust the enterprise CA. You can either generate the needed certificates and import them onto the firewall, or generate a certificate signing request (CSR) on the firewall and send it to the enterprise CA for signing. The benefit of this method is that the private key does not leave the firewall. An enterprise CA can also issue a signing certificate, which the firewall uses to automatically generate certificates (for example, for GlobalProtect large-scale VPN or sites requiring SSL/TLS decryption). See Import a Certificate and Private Key.

- **Generate self-signed certificates**—You can Create a Self-Signed Root CA Certificate on the firewall and use it to automatically issue certificates for other firewall applications. Note that if you use this method to generate certificates for an application that requires an end client to trust the certificate, end users will see a certificate error because the root CA certificate is not in their trusted root certificate store. To prevent this, deploy the self-signed root CA certificate to all end user systems. You can deploy the certificates manually or use a centralized deployment method such as an Active Directory Group Policy Object (GPO).

 © Palo Alto Networks, Inc.

# Set Up Verification for Certificate Revocation Status

To verify the revocation status of certificates, the firewall uses Online Certificate Status Protocol (OCSP) and/or certificate revocation lists (CRLs). For details on these methods, see Certificate Revocation If you configure both methods, the firewall first tries OCSP and only falls back to the CRL method if the OCSP responder is unavailable. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall to function as the OCSP responder.

The following topics describe how to configure the firewall to verify certificate revocation status:

- Configure an OCSP Responder
- Configure Revocation Status Verification of Certificates Used for User/Device Authentication
- Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

## Configure an OCSP Responder

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of certificates, you must configure the firewall to access an OCSP responder (server). The entity that manages the OCSP responder can be a third-party certificate authority (CA) or, if your enterprise has its own public key infrastructure (PKI), the firewall itself. For details on OCSP, see Certificate Revocation

| **Configure an OCSP Responder** | | |
|---|---|---|
| Step 1 | Define an OCSP responder. | 1. Select **Device > Certificate Management > OCSP Responder** and click **Add**.<br>2. Enter a **Name** to identify the responder (up to 31 characters). The name is case-sensitive. It must be unique and use only letters, numbers, spaces, hyphens, and underscores.<br>3. If the device has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.<br>4. In the **Host Name** field, enter the host name (recommended) or IP address of the OCSP responder. From this value, PAN-OS automatically derives a URL and adds it to the certificate being verified.<br>If you configure the firewall itself as an OCSP responder, the host name must resolve to an IP address in the interface that the firewall uses for OCSP services (specified in Step 3).<br>5. Click **OK**. |
| Step 2 | Enable OCSP communication on the firewall. | 1. Select **Device > Setup > Management**.<br>2. In the Management Interface Settings section, edit to select the **HTTP OCSP** check box, then click **OK**. |

| Configure an OCSP Responder | | |
|---|---|---|
| Step 3 | Optionally, to configure the firewall itself as an OCSP responder, add an Interface Management Profile to the interface used for OCSP services. | 1. Select **Network > Network Profiles > Interface Mgmt**.<br>2. Click **Add** to create a new profile or click the name of an existing profile.<br>3. Select the **HTTP OCSP** check box and click **OK**.<br>4. Select **Network > Interfaces** and click the name of the interface that the firewall will use for OCSP services. The OCSP **Host Name** specified in Step 1 must resolve to an IP address in this interface.<br>5. Select **Advanced > Other info** and select the Interface Management Profile you configured.<br>6. Click **OK** and **Commit**. |

## Configure Revocation Status Verification of Certificates Used for User/Device Authentication

The firewall uses certificates to authenticate users and devices for such applications as Captive Portal, GlobalProtect, site-to-site IPSec VPN, and web interface access to Palo Alto Networks devices. To improve security, it is a best practice to configure the firewall to verify the revocation status of certificates that it uses for device/user authentication.

| Configure Revocation Status Verification of Certificates Used for User/Device Authentication | | |
|---|---|---|
| Step 1 | Configure a Certificate Profile for each application. | Assign one or more root CA certificates to the profile and select how the firewall verifies certificate revocation status. The common name (FQDN or IP address) of a certificate must match an interface to which you apply the profile in Step 2.<br>For details on the certificates that various applications use, see Keys and Certificates |
| Step 2 | Assign the certificate profiles to the relevant applications. | The steps to assign a certificate profile depend on the application that requires it. |

## Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

The firewall decrypts inbound and outbound SSL/TLS traffic to apply security rules and rules, then re-encrypts the traffic before forwarding it. (For details, see SSL Inbound Inspection and SSL Forward Proxy.) You can configure the firewall to verify the revocation status of certificates used for decryption as follows.

Enabling revocation status verification for SSL/TLS decryption certificates will add time to the process of establishing the session. The first attempt to access a site might fail if the verification does not finish before the session times out. For these reasons, verification is disabled by default.

 © Palo Alto Networks, Inc.

| Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption | | |
|---|---|---|
| Step 1 | Define the service-specific timeout intervals for revocation status requests. | 1. Select **Device > Setup > Session** and, in the Session Features section, select **Decryption Certificate Revocation Settings**.<br>2. Perform one or both of the following steps, depending on whether the firewall will use Online Certificate Status Protocol (OCSP) or the Certificate Revocation List (CRL) method to verify the revocation status of certificates. If the firewall will use both, it first tries OCSP; if the OCSP responder is unavailable, the firewall then tries the CRL method.<br>• In the CRL section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the CRL service.<br>• In the OCSP section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the OCSP responder.<br>Depending on the **Certificate Status Timeout** value you specify in Step 2, the firewall might register a timeout before either or both of the **Receive Timeout** intervals pass. |
| Step 2 | Define the total timeout interval for revocation status requests. | Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies the session-blocking logic you optionally define in Step 3. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:<br>• If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.<br>• If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.<br>• If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value. |
| Step 3 | Define the blocking behavior for *unknown* certificate status or a revocation status request timeout. | If you want the firewall to block SSL/TLS sessions when the OCSP or CRL service returns a certificate revocation status of *unknown*, select the **Block Session With Unknown Certificate Status** check box. Otherwise, the firewall proceeds with the session.<br>If you want the firewall to block SSL/TLS sessions after it registers a request timeout, select the **Block Session On Certificate Status Check Timeout** check box. Otherwise, the firewall proceeds with the session. |
| Step 4 | Save and apply your entries. | Click **OK** and **Commit**. |

# Configure the Master Key

Every firewall has a default master key that encrypts private keys and other secrets (such as passwords and shared keys). The private keys authenticate users when they access administrative interfaces on the firewall. As a best practice to safeguard the keys, configure the master key on each firewall to be unique and periodically change it. For added security, use a *wrapping key* stored on a hardware security module (HSM) to encrypt the master key. For details, see Encrypt a Master Key Using an HSM.

In a high availability (HA) configuration, ensure both devices in the pair use the same master key to encrypt private keys and certificates. If the master keys differ, HA configuration synchronization will not work properly.

When you export a firewall configuration, the master key encrypts the passwords of users managed on external servers. For locally managed users, the firewall hashes the passwords but the master key does not encrypt them.

| **Configure a Master Key** | |
|---|---|
| Step 1 | Select **Device > Master Key and Diagnostics** and, in the Master Key section, click the Edit icon. |
| Step 2 | Enter the **Current Master Key** if one exists. |
| Step 3 | Define a new **New Master Key** and then **Confirm New Master Key**. The key must contain exactly 16 characters. |
| Step 4 | (Optional) To specify the master key **Life Time**, enter the number of **Days** and/or **Hours** after which the key will expire. If you set a life time, create a new master key before the old key expires. |
| Step 5 | (Optional) If you set a key life time, enter a **Time for Reminder** that specifies the number of **Days** and **Hours** preceding master key expiration when the firewall emails you a reminder. |
| Step 6 | (Optional) Select whether to use an **HSM** to encrypt the master key. For details, see Encrypt a Master Key Using an HSM. |
| Step 7 | Click **OK** and **Commit**. |


© Palo Alto Networks, Inc.

# Obtain Certificates


- Create a Self-Signed Root CA Certificate
- Generate a Certificate on the Device
- Import a Certificate and Private Key
- Obtain a Certificate from an External CA


© Palo Alto Networks, Inc.

## Create a Self-Signed Root CA Certificate

A self-signed root certificate authority (CA) certificate is the top-most certificate in a certificate chain. A firewall can use this certificate to automatically issue certificates for other uses. For example, the firewall issues certificates for SSL/TLS decryption and for satellite devices in a GlobalProtect large-scale VPN.

When establishing a secure connection with the firewall, the remote client must trust the root CA that issued the certificate. Otherwise, the client browser will display a warning that the certificate is invalid and might (depending on security settings) block the connection. To prevent this, after generating the self-signed root CA certificate, import it into the client systems.

On a Palo Alto Networks device, you can generate self-signed certificates only if they are CA certificates.

| Generate a Self-signed Root CA Certificate | |
|---|---|
| Step 1 | Select **Device > Certificate Management > Certificates > Device Certificates**. |
| Step 2 | If the device has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate. |
| Step 3 | Click **Generate**. |
| Step 4 | Enter a **Certificate Name**, such as *GlobalProtect_CA*. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores. |
| Step 5 | In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate. |
| Step 6 | If the device has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box. |
| Step 7 | Leave the **Signed By** field blank to designate the certificate as self-signed. |
| Step 8 | (Required) Select the **Certificate Authority** check box. |
| Step 9 | Leave the **OCSP Responder** field blank; revocation status verification doesn't apply to root CA certificates. |
| Step 10 | Click **Generate** and **Commit**. |

 © Palo Alto Networks, Inc.

## Generate a Certificate on the Device

Palo Alto Networks devices use certificates to authenticate clients, servers, users, and devices in several applications, including SSL/TLS decryption, Captive Portal, GlobalProtect, site-to-site IPSec VPN, and device web interface access. Generate certificates for each usage: for details, see Keys and Certificates.

To generate a certificate, you must first Create a Self-Signed Root CA Certificate or import one (Import a Certificate and Private Key) to sign it. To use Online Certificate Status Protocol (OCSP) for verifying certificate revocation status, Configure an OCSP Responder before generating the certificate.

| Generate a Certificate on the Device | |
|---|---|
| Step 1 | Select **Device > Certificate Management > Certificates > Device Certificates**. |
| Step 2 | If the device has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate. |
| Step 3 | Click **Generate**. |
| Step 4 | Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores. |
| Step 5 | In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate. |
| Step 6 | If the device has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box. |
| Step 7 | In the **Signed By** field, select the root CA certificate that will issue the certificate. |
| Step 8 | (Optional) Select an **OCSP Responder**. |
| Step 9 | For the key generation **Algorithm**, select **RSA** (default) or **Elliptical Curve DSA** (ECDSA). ECDSA is recommended for client browsers and operating systems that support it.<br> Firewalls that run PAN-OS 6.1 and earlier releases will delete any ECDSA certificates that you push from Panorama™, and any RSA certificates signed by an ECDSA certificate authority (CA) will be invalid on those firewalls. |
| Step 10 | Select the **Number of Bits** to define the certificate key length. Higher numbers are more secure but require more processing time. |
| Step 11 | Select the **Digest** algorithm. From most to least secure, the options are: **sha512**, **sha384**, **sha256** (default), **sha1**, and **md5**. |
| Step 12 | For the **Expiration**, enter the number of days (default is 365) for which the certificate is valid. |
| Step 13 | (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.<br> If you add a **Host Name** (DNS name) attribute, it is a best practice for it to match the **Common Name**. The host name populates the Subject Alternative Name field of the certificate. |
| Step 14 | Click **Generate** and, in the Device Certificates page, click the certificate Name.<br> Regardless of the time zone on the firewall, it always displays the corresponding Greenwich Mean Time (GMT) for certificate validity and expiration dates/times. |

| Generate a Certificate on the Device (Continued) | |
|---|---|
| Step 15 | Select the check boxes that correspond to the intended use of the certificate on the firewall.<br>For example, if the firewall will use this certificate to secure forwarding of syslogs to an external syslog server, select the **Certificate for Secure Syslog** check box. |
| Step 16 | Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

## Import a Certificate and Private Key

If your enterprise has its own public key infrastructure (PKI), you can import a certificate and private key into the firewall from your enterprise certificate authority (CA). Enterprise CA certificates (unlike most certificates purchased from a trusted, third-party CA) can automatically issue CA certificates for applications such as SSL/TLS decryption or large-scale VPN.

On a Palo Alto Networks device, you can import self-signed certificates only if they are CA certificates.

Instead of importing a self-signed root CA certificate into all the client systems, it is a best practice to import a certificate from the enterprise CA because the clients will already have a trust relationship with the enterprise CA, which simplifies the deployment.

If the certificate you will import is part of a certificate chain, it is a best practice to import the entire chain.

| Import a Certificate and Private Key | |
|---|---|
| Step 1 | From the enterprise CA, export the certificate and private key that the firewall will use for authentication.<br><br>When exporting a private key, you must enter a passphrase to encrypt the key for transport. Ensure the management system can access the certificate and key files. When importing the key onto the firewall, you must enter the same passphrase to decrypt it. |
| Step 2 | Select **Device > Certificate Management > Certificates > Device Certificates**. |
| Step 3 | If the device has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate. |
| Step 4 | Click **Import** and enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores. |
| Step 5 | To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the device supports multiple virtual systems. |
| Step 6 | Enter the path and name of the **Certificate File** received from the CA, or **Browse** to find the file. |
| Step 7 | Select a **File Format**:<br>• **Encrypted Private Key and Certificate (PKCS12)**—This is the default and most common format, in which the key and certificate are in a single container (**Certificate File**). If a hardware security module (HSM) will store the private key for this certificate, select the **Private key resides on Hardware Security Module** check box.<br>• **Base64 Encoded Certificate (PEM)**—You must import the key separately from the certificate. If a hardware security module (HSM) stores the private key for this certificate, select the **Private key resides on Hardware Security Module** check box and skip Step 8. Otherwise, select the **Import Private Key** check box, enter the **Key File** or **Browse** to it, then perform Step 8. |
| Step 8 | Enter and re-enter (confirm) the **Passphrase** used to encrypt the private key. |
| Step 9 | Click **OK**. The Device Certificates page displays the imported certificate. |

## Obtain a Certificate from an External CA

The advantage of obtaining a certificate from an external certificate authority (CA) is that the private key does not leave the firewall. To obtain a certificate from an external CA, generate a certificate signing request (CSR) and submit it to the CA. After the CA issues a certificate with the specified attributes, import it onto the firewall. The CA can be a well-known, public CA or an enterprise CA.

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of the certificate, Configure an OCSP Responder *before* generating the CSR.

| **Obtain a Certificate from an External CA** | |
|---|---|
| Step 1 Request the certificate from an external CA. | 1. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. If the device has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.<br>3. Click **Generate**.<br>4. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.<br>5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.<br>6. If the device has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box.<br>7. In the **Signed By** field, select **External Authority (CSR)**.<br>8. If applicable, select an **OCSP Responder**.<br>9. (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.<br>If you add a **Host Name** attribute, it is a best practice for it to match the **Common Name** (this is mandatory for GlobalProtect). The host name populates the Subject Alternative Name field of the certificate.<br>10. Click **Generate**. The Device Certificates tab displays the CSR with a Status of *pending*. |
| Step 2 Submit the CSR to the CA. | 1. Select the CSR and click **Export** to save the .csr file to a local computer.<br>2. Upload the .csr file to the CA. |


© Palo Alto Networks, Inc.

| Obtain a Certificate from an External CA | |
|---|---|
| Step 3 Import the certificate. | 1. After the CA sends a signed certificate in response to the CSR, return to the **Device Certificates** tab and click **Import**.<br>2. Enter the **Certificate Name** used to generate the CSR in Step 1-4.<br>3. Enter the path and name of the PEM **Certificate File** that the CA sent, or **Browse** to it.<br>4. Click **OK**. The **Device Certificates** tab displays the certificate with a Status of *valid*. |
| Step 4 Configure the certificate. | 1. Click the certificate **Name**.<br>2. Select the check boxes that correspond to the intended use of the certificate on the firewall. For example, if the firewall will use this certificate to secure forwarding of syslogs to an external syslog server, select the **Certificate for Secure Syslog** check box.<br>3. Click **OK** and **Commit**. |

# Export a Certificate and Private Key

Palo Alto Networks recommends that you use your enterprise public key infrastructure (PKI) to distribute a certificate and private key in your organization. However, if necessary, you can also export a certificate and private key from the firewall or Panorama. You can use an exported certificate and private key in the following cases:

- Administrator authentication to the device web interface
- GlobalProtect agent/app authentication to portals and gateways
- SSL Forward Proxy decryption
- Certificate signing requests (CSRs)

| Export a Certificate and Private Key | |
|---|---|
| Step 1 | Select **Device > Certificate Management > Certificates > Device Certificates**. |
| Step 2 | If the device has more than one virtual system (vsys), select a **Location** (a specific vsys or **Shared**) for the certificate. |
| Step 3 | Select the certificate, click **Export**, and select a **File Format**:<br>• **Base64 Encoded Certificate (PEM)**—This is the default format. It is the most common and has the broadest support on the Internet. If you want the exported file to include the private key, select the **Export Private Key** check box.<br>• **Encrypted Private Key and Certificate (PKCS12)**—This format is more secure than PEM but is not as common or as broadly supported. The exported file will automatically include the private key.<br>• **Binary Encoded Certificate (DER)**—More operating system types support this format than the others. You can export only the certificate, not the key: ignore the **Export Private Key** check box and passphrase fields. |
| Step 4 | Enter a **Passphrase** and **Confirm Passphrase** to encrypt the private key if the **File Format** is PKCS12 or if it is PEM and you selected the **Export Private Key** check box. You will use this passphrase when importing the certificate and key into client systems. |
| Step 5 | Click **OK** and save the certificate/key file to your computer. |


© Palo Alto Networks, Inc.

# Configure a Certificate Profile

Certificate profiles define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPSec VPN, Mobile Security Manager, and web interface access to Palo Alto Networks devices. The profiles specify which certificates to use, how to verify certificate revocation status, and how that status constrains access. Configure a certificate profile for each application.

It is a best practice to enable Online Certificate Status Protocol (OCSP) and/or Certificate Revocation List (CRL) status verification for certificate profiles. For details on these methods, see Certificate Revocation.

| Configure a Certificate Profile | | |
|---|---|---|
| Step 1 | Obtain the certificate authority (CA) certificates you will assign. | Perform one of the following steps to obtain the CA certificates you will assign to the profile. You must assign at least one.<br>• Generate a Certificate on the Device.<br>• Export a certificate from your enterprise CA and then import it onto the firewall (see Step 3). |
| Step 2 | Identify the certificate profile. | 1. Select **Device > Certificate Management > Certificates Profile** and click **Add**.<br>2. Enter a **Name** to identify the profile. The name is case-sensitive, must be unique and can use up to 31 characters that include only letters, numbers, spaces, hyphens, and underscores.<br>3. If the device has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate. |
| Step 3 | Assign one or more certificates. | Perform the following steps for each CA certificate:<br>1. In the CA Certificates table, click **Add**.<br>2. Select a **CA Certificate**. Alternatively, to import a certificate, click **Import**, enter a **Certificate Name**, **Browse** to the **Certificate File** you exported from your enterprise CA, and click **OK**.<br>3. Optionally, if the firewall uses OCSP to verify certificate revocation status, configure the following fields to override the default behavior. For most deployments, these fields do not apply.<br>• By default, the firewall uses the OCSP responder URL that you set in the procedure Configure an OCSP Responder. To override that setting, enter a **Default OCSP URL** (starting with *http://* or *https://*).<br>• By default, the firewall uses the certificate selected in the **CA Certificate** field to validate OCSP responses. To use a different certificate for validation, select it in the **OCSP Verify CA Certificate** field.<br>4. Click **OK**. The CA Certificates table displays the assigned certificate. |

| Configure a Certificate Profile | |
|---|---|
| Step 4 Define the methods for verifying certificate revocation status and the associated blocking behavior. | 1. Select **Use CRL** and/or **Use OCSP**. If you select both, the firewall first tries OCSP and falls back to the CRL method only if the OCSP responder is unavailable.<br>2. Depending on the verification method, enter the **CRL Receive Timeout** and/or **OCSP Receive Timeout**. These are the intervals (1-60 seconds) after which the firewall stops waiting for a response from the CRL/OCSP service.<br>3. Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies any session-blocking logic you define. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:<br>• If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.<br>• If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.<br>• If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value.<br>4. If you want the firewall to block sessions when the OCSP or CRL service returns a certificate revocation status of *unknown*, select the **Block session if certificate status is unknown** check box. Otherwise, the firewall proceeds with the session.<br>5. If you want the firewall to block sessions after it registers an OCSP or CRL request timeout, select the **Block session if certificate status cannot be retrieved within timeout** check box. Otherwise, the firewall proceeds with the session. |
| Step 5 Save and apply your entries. | Click **OK** and **Commit**. |


© Palo Alto Networks, Inc.

# Configure an SSL/TLS Service Profile

SSL/TLS service profiles specify a certificate and the allowed protocol versions for the SSL/TLS services of Palo Alto Networks devices. The devices use SSL/TLS for Captive Portal, GlobalProtect portals and gateways, inbound traffic on the management (MGT) interface, the URL Admin Override feature, and the User-ID™ syslog listening service. By defining the protocol versions, you can use a profile to restrict the cipher suites that are available for securing communication with the clients requesting the services. This improves network security by enabling devices to avoid SSL/TLS versions that have known weaknesses: if a service request involves a protocol version that is outside the specified range, the device downgrades or upgrades the connection to a supported version.

| Configure an SSL/TLS Service Profile | |
|---|---|
| Step 1 | For each desired service, generate or import a certificate on the firewall (see Obtain Certificates).<br> Use only signed certificates, not certificate authority (CA) certificates, for SSL/TLS services. |
| Step 2 | Select **Device > Certificate Management > SSL/TLS Service Profile**. |
| Step 3 | If the device has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available. |
| Step 4 | Click **Add** and enter a **Name** to identify the profile. |
| Step 5 | Select the **Certificate** you just obtained. |
| Step 6 | Define the range of protocols that the service can use:<br>• For the **Min Version**, select the earliest allowed TLS version: **TLSv1.0** (default), **TLSv1.1**, or **TLSv1.2**.<br>• For the **Max Version**, select the latest allowed TLS version: **TLSv1.0**, **TLSv1.1**, **TLSv1.2**, or **Max** (latest available version). The default is **Max**. |
| Step 7 | Click **OK** and **Commit**. |

# Configure the Key Size for SSL Forward Proxy Server Certificates

When responding to a client in an SSL Forward Proxy session, the firewall creates a copy of the certificate that the destination server presents and uses the copy to establish a connection with the client. By default, the firewall generates certificates with the same key size as the certificate that the destination server presented. However, you can change the key size for the firewall-generated certificate as follows:

**Configure the Key Size for SSL Forward Proxy Server Certificates**

Step 1 Select **Device > Setup > Session** and, in the Decryption Settings section, click **SSL Forward Proxy Settings**.

Step 2 Select a **Key Size**:

- **Defined by destination host**—The firewall determines the key size for the certificates it generates to establish SSL proxy sessions with clients based on the key size of the destination server certificate. If the destination server uses a 1,024-bit RSA key, the firewall generates a certificate with that key size and an SHA-1 hashing algorithm. If the destination server uses a key size larger than 1,024 bits (for example, 2,048 bits or 4,096 bits), the firewall generates a certificate that uses a 2,048-bit RSA key and SHA-256 algorithm. This is the default setting.
- **1024-bit RSA**—The firewall generates certificates that use a 1,024-bit RSA key and SHA-1 hashing algorithm regardless of the key size of the destination server certificates. As of December 31, 2013, public certificate authorities (CAs) and popular browsers have limited support for X.509 certificates that use keys of fewer than 2,048 bits. In the future, depending on security settings, when presented with such keys the browser might warn the user or block the SSL/TLS session entirely.
- **2048-bit RSA**—The firewall generates certificates that use a 2,048-bit RSA key and SHA-256 hashing algorithm regardless of the key size of the destination server certificates. Public CAs and popular browsers support 2,048-bit keys, which provide better security than the 1,024-bit keys.

Changing the key size setting clears the current certificate cache.

Step 3 Click **OK** and **Commit**.

 © Palo Alto Networks, Inc.

# Revoke and Renew Certificates

▲ Revoke a Certificate
▲ Renew a Certificate

## Revoke a Certificate

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority (CA) that issued the certificate must revoke it. The following task describes how to revoke a certificate for which the firewall is the CA.

| **Revoke a Certificate** | |
|---|---|
| Step 1 | Select **Device > Certificate Management > Certificates > Device Certificates**. |
| Step 2 | If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select the virtual system to which the certificate belongs. |
| Step 3 | Select the certificate to revoke. |
| Step 4 | Click **Revoke**. PAN-OS immediately sets the status of the certificate to revoked and adds the serial number to the Online Certificate Status Protocol (OCSP) responder cache or certificate revocation list (CRL). You need not perform a commit. |

## Renew a Certificate

If a certificate expires, or soon will, you can reset the validity period. If an external certificate authority (CA) signed the certificate and the firewall uses the Online Certificate Status Protocol (OCSP) to verify certificate revocation status, the firewall uses the OCSP responder information to update the certificate status (see Configure an OCSP Responder). If the firewall is the CA that issued the certificate, the firewall replaces it with a new certificate that has a different serial number but the same attributes as the old certificate.

| **Renew a Certificate** | |
|---|---|
| Step 1 | Select **Device > Certificate Management > Certificates > Device Certificates**. |
| Step 2 | If the device has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate. |
| Step 3 | Select a certificate to renew and click **Renew**. |
| Step 4 | Enter a **New Expiration Interval** (in days). |
| Step 5 | Click **OK** and **Commit**. |

© Palo Alto Networks, Inc.

# Secure Keys with a Hardware Security Module

A hardware security module (HSM) is a physical device that manages digital keys. An HSM provides secure storage and generation of digital keys. It provides both logical and physical protection of these materials from non-authorized use and potential adversaries.

HSM clients integrated with Palo Alto Networks devices enable enhanced security for the private keys used in SSL/TLS decryption (both SSL forward proxy and SSL inbound inspection). In addition, you can use the HSM to encrypt device master keys.

The following topics describe how to integrate an HSM with your Palo Alto Networks devices:

- Set up Connectivity with an HSM
- Encrypt a Master Key Using an HSM
- Store Private Keys on an HSM
- Manage the HSM Deployment

 © Palo Alto Networks, Inc.

# Set up Connectivity with an HSM

HSM clients are integrated with PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7000 Series, and VM-Series firewalls and on Panorama (virtual appliance and M-Series appliance) for use with the following HSMs:

- SafeNet Network 5.2.1 or later
- Thales Nshield Connect 11.62 or later

The HSM server version must be compatible with these client versions. Refer to the HSM vendor documentation for the client-server version compatibility matrix.

The following topics describe how to set up connectivity between the Palo Alto Networks device and one of the supported HSMs:

▲ Set Up Connectivity with a SafeNet Network HSM

▲ Set Up Connectivity with a Thales Nshield Connect HSM

## Set Up Connectivity with a SafeNet Network HSM

To set up connectivity between the Palo Alto Networks device and a SafeNet Network HSM, you must specify the address of the HSM server and the password for connecting to it in the firewall configuration. In addition, you must register the firewall with the HSM server. Prior to beginning the configuration, make sure you have created a partition for the Palo Alto Networks devices on the HSM server.

HSM configuration is not synced between high availability firewall peers. Consequently, you must configure the HSM module separately on each of the peers.

In Active-Passive HA deployments, you must manually perform one failover to configure and authenticate each HA peer individually to the HSM. After this manual failover has been performed, user interaction is not required for the failover function.

| **Set up a Connectivity with a SafeNet Network HSM** | |
|---|---|
| Step 1 Configure the firewall to communicate with the SafeNet Network HSM. | 1. Log in to the firewall web interface and select **Device > Setup > HSM**.<br>2. Edit the Hardware Security Module Provider section and select **Safenet Luna SA** (SafeNet Network) as the **Provider Configured**.<br>3. Click **Add** and enter a **Module Name**. This can be any ASCII string up to 31 characters in length.<br>4. Enter the IPv4 address of the HSM module as the **Server Address**.<br>If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices.<br>5. (Optional) If configuring a high availability HSM configuration, select the **High Availability** check box and add the following: a value for **Auto Recovery Retry** and a **High Availability Group Name**.<br>If two HSM servers are configured, you should configure high availability. Otherwise the second HSM server is not used.<br>6. Click **OK** and **Commit.** |

| Set up a Connectivity with a SafeNet Network HSM (Continued) | | |
|---|---|---|
| Step 2 | (Optional) Configure a service route to enable the firewall to connect to the HSM.<br><br>By default, the firewall uses the Management Interface to communicate with the HSM. To use a different interface, you must configure a service route. | 1. Select **Device > Setup > Services.**<br>2. Select **Service Route Configuration** from the Services Features area.<br>3. Select **Customize** from the Service Route Configuration area.<br>4. Select the **IPv4** tab.<br>5. Select **HSM** from the **Service** column.<br>6. Select an interface to use for HSM from the **Source Interface** drop-down.<br><br>If you select a dataplane connected port for HSM, issuing the `clear session all` CLI command will clear all existing HSM sessions, causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.<br><br>7. Click **OK** and **Commit.** |
| Step 3 | Configure the firewall to authenticate to the HSM. | 1. Select **Device > Setup > HSM**.<br>2. Select **Setup Hardware Security Module** in the Hardware Security Operations area.<br>3. Select the HSM **Server Name** from the drop-down.<br>4. Enter the **Administrator Password** to authenticate the firewall to the HSM.<br>5. Click **OK.**<br><br>The firewall attempts to perform an authentication with the HSM and displays a status message.<br><br>6. Click **OK.** |
| Step 4 | Register the firewall (the HSM client) with the HSM and assign it to a partition on the HSM.<br><br>If the HSM already has a firewall with the same `<cl-name>` registered, you must remove the duplicate registration using the following command before registration will succeed:<br><br>`client delete -client <cl-name>`<br><br>where `<cl-name>` is the name of the client (firewall) registration you want to delete. | 1. Log in to the HSM from a remote system.<br>2. Register the firewall using the following command:<br>`client register -c <cl-name> -ip <fw-ip-addr>`<br>where `<cl-name>` is a name that you assign to the firewall for use on the HSM and `<fw-ip-addr>` is the IP address of the firewall that is being configured as an HSM client.<br>3. Assign a partition to the firewall using the following command:<br>`client assignpartition -c <cl-name> -p <partition-name>`<br>where `<cl-name>` is the name assigned to the firewall in the `client register` command and `<partition-name>` is the name of a previously configured partition that you want to assign to the firewall. |

 © Palo Alto Networks, Inc.

| Set up a Connectivity with a SafeNet Network HSM (Continued) | | |
|---|---|---|
| Step 5 | Configure the firewall to connect to the HSM partition. | 1. Select **Device > Setup > HSM**.<br>2. Click the **Refresh** icon.<br>3. Select the **Setup HSM Partition** in the Hardware Security Operations area.<br>4. Enter the **Partition Password** to authenticate the firewall to the partition on the HSM.<br>5. Click **OK**. |
| Step 6 | (Optional) Configure an additional HSM for high availability (HA). | 1. Follow Step 1 through Step 5 to add an additional HSM for high availability (HA).<br>This process adds a new HSM to the existing HA group.<br>2. If you remove an HSM from your configuration, repeat Step 5.<br>This will remove the deleted HSM from the HA group. |
| Step 7 | Verify connectivity with the HSM. | 1. Select **Device > Setup > HSM**.<br>2. Check the **Status** of the HSM connection:<br>Green—HSM is authenticated and connected.<br>Red—HSM was not authenticated or network connectivity to the HSM is down.<br>3. View the following columns in Hardware Security Module Status area to determine authentication status:<br>**Serial Number**—The serial number of the HSM partition if the HSM was successfully authenticated.<br>**Partition—**The partition name on the HSM that was assigned on the firewall.<br>**Module State**—The current operating state of the HSM. It always has the value Authenticated if the HSM is displayed in this table. |

## Set Up Connectivity with a Thales Nshield Connect HSM

The following workflow describes how to configure the firewall to communicate with a Thales Nshield Connect HSM. This configuration requires that you set up a remote filesystem (RFS) to use as a *hub* to sync key data for all firewalls in your organization that are using the HSM.

HSM configuration is not synced between high availability firewall peers. Consequently, you must configure the HSM module separately on each of the peers.

If the high availability firewall configuration is in Active-Passive mode, you must manually perform one failover to configure and authenticate each HA peer individually to the HSM. After this manual failover has been performed, user interaction is not required for the failover function.

| Set up Connectivity with a Thales Nshield Connect HSM | | |
|---|---|---|
| Step 1 | Configure the Thales Nshield Connect server as the firewall's HSM provider. | 1. From the firewall web interface, select **Device > Setup > HSM** and edit the Hardware Security Module Provider section.<br>2. Select **Thales Nshield Connect** as the **Provider Configured**.<br>3. Click **Add** and enter a **Module Name**. This can be any ASCII string up to 31 characters in length.<br>4. Enter the IPv4 address as the **Server Address** of the HSM module.<br>If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices.<br>5. Enter the IPv4 address of the **Remote Filesystem Address**.<br>6. Click **OK** and **Commit**. |
| Step 2 | (Optional) Configure a service route to enable the firewall to connect to the HSM.<br>By default, the firewall uses the Management Interface to communicate with the HSM. To use a different interface, you must configure a service route. | 1. Select **Device > Setup > Services.**<br>2. Select **Service Route Configuration** from the Services Features area.<br>3. Select **Customize** from the Service Route Configuration area.<br>4. Select the **IPv4** tab.<br>5. Select **HSM** from the **Service** column.<br>6. Select an interface to use for HSM from the **Source Interface** drop-down.<br>If you select a dataplane connected port for HSM, issuing the `clear session all` CLI command will clear all existing HSM sessions, causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.<br>7. Click **OK** and **Commit.** |
| Step 3 | Register the firewall (the HSM client) with the HSM server.<br>This step briefly describes the procedure for using the front panel interface of the Thales Nshield Connect HSM. For more details, consult the Thales documentation. | 1. Log in to the front panel display of the Thales Nshield Connect HSM unit.<br>2. On the unit front panel, use the right-hand navigation button to select **System > System configuration > Client config > New client**.<br>3. Enter the IP address of the firewall.<br>4. Select **System > System configuration > Client config > Remote file system** and enter the IP address of the client computer where you set up the remote file system. |

 © Palo Alto Networks, Inc.

| Set up Connectivity with a Thales Nshield Connect HSM (Continued) | | |
|---|---|---|
| Step 4 | Set up the remote filesystem to accept connections from the firewall. | 1. Log in to the remote filesystem (RFS) from a Linux client.<br>2. Obtain the electronic serial number (ESN) and the hash of the $K_{NETI}$ key. The $K_{NETI}$ key authenticates the module to clients:<br>`anonkneti <ip-address>`<br>where `<ip-address>` is the IP address of the HSM.<br>The following is an example:<br>`anonkneti 192.0.2.1`<br>`B1E2-2D4C-E6A2 5a2e5107e70d525615a903f6391ad72b1c03352c`<br>In this example, **`B1E2-2D4C-E6A2`** is the ESM and 5a2e5107e70d525615a903f6391ad72b1c03352c is the hash of the $K_{NETI}$ key.<br>3. Use the following command from a superuser account to perform the remote filesystem setup:<br>`rfs-setup --force <ip-address> <ESN> <hash-Kneti-key>`<br>where `<ip-address>` is the IP address of the HSM,<br>`<ESN>` is the electronic serial number (ESN) and<br>`<hash-Kneti-key>` is the hash of the KNETI key.<br>The following example uses the values obtained in this procedure:<br>`rfs-setup --force <192.0.2.1> <B1E2-2D4C-E6A2>`<br>`<5a2e5107e70d525615a903f6391ad72b1c03352c>`<br>4. Use the following command to permit client submit on the Remote Filesystem:<br>`rfs-setup --gang-client --write-noauth <FW-IPaddress>`<br>where `<FW-IPaddress>` is the IP address of the firewall. |
| Step 5 | Configure the firewall to authenticate to the HSM. | 1. From the firewall web interface, select **Device > Setup > HSM**.<br>2. Select **Setup Hardware Security Module** in the Hardware Security Operations area.<br>3. Click **OK.**<br>The firewall attempts to perform an authentication with the HSM and displays a status message.<br>4. Click **OK.** |
| Step 6 | Synchronize the firewall with the remote filesystem. | 1. Select the **Device > Setup > HSM**.<br>2. Select **Synchronize with Remote Filesystem** in the Hardware Security Operations section. |

| Set up Connectivity with a Thales Nshield Connect HSM (Continued) | | |
|---|---|---|
| Step 7 | Verify that the firewall can connect to the HSM. | 1. Select **Device > Setup > HSM**.<br>2. Check the Status indicator to verify that the firewall is connected to the HSM:<br>Green—HSM is authenticated and connected.<br>Red—HSM was not authenticated or network connectivity to the HSM is down.<br>3. View the following columns in Hardware Security Module Status section to determine authentication status.<br>**Name**: The name of the HSM attempting to be authenticated.<br>**IP address:** The IP address of the HSM that was assigned on the firewall.<br>**Module State**: The current operating state of the HSM: **Authenticated** or **Not Authenticated**. |


© Palo Alto Networks, Inc.

# Encrypt a Master Key Using an HSM

A master key is configured on a Palo Alto Networks firewall to encrypt all private keys and passwords. If you have security requirements to store your private keys in a secure location, you can encrypt the master key using an encryption key that is stored on an HSM. The firewall then requests the HSM to decrypt the master key whenever it is required to decrypt a password or private key on the firewall. Typically, the HSM is located in a highly secure location that is separate from the firewall for greater security.

The HSM encrypts the master key using a wrapping key. To maintain security, this encryption key must occasionally be changed. For this reason, a command is provided on the firewall to rotate the wrapping key which changes the master key encryption. The frequency of this wrapping key rotation depends on your application.

Master key encryption using an HSM is not supported on firewalls configured in FIPS or CC mode.

The following topics describe how to encrypt the master key initially and how to refresh the master key encryption:

- Encrypt the Master Key
- Refresh the Master Key Encryption

## Encrypt the Master Key

If you have not previously encrypted the master key on a device, use the following procedure to encrypt it. Use this procedure for first time encryption of a key, or if you define a new master key and you want to encrypt it. If you want to refresh the encryption on a previously encrypted key, see Refresh the Master Key Encryption.

| Encrypt a Master Key Using an HSM | |
|---|---|
| Step 1 | Select **Device > Master Key and Diagnostics.** |
| Step 2 | Specify the key that is currently used to encrypt all of the private keys and passwords on the firewall in the **Master Key** field. |
| Step 3 | If changing the master key, enter the new master key and confirm. |
| Step 4 | Select the **HSM** check box.<br>**Life Time:** The number of days and hours after which the master key expires (range 1-730 days).<br>**Time for Reminder:** The number of days and hours before expiration when the user is notified of the impending expiration (range 1-365 days). |
| Step 5 | Click **OK**. |

## Refresh the Master Key Encryption

As a best practice, refresh the master key encryption on a regular basis by rotating the master key wrapping key on the HSM. This command is the same for both the SafeNet Network and Thales Nshield Connect HSMs.

**Refresh the Master Key Encryption**

Step 1 Use the following CLI command to rotate the wrapping key for the master key on an HSM:

```
> request hsm mkey-wrapping-key-rotation
```

If the master key is encrypted on the HSM, the CLI command will generate a new wrapping key on the HSM and encrypt the master key with the new wrapping key.

If the master key is not encrypted on the HSM, the CLI command will generate new wrapping key on the HSM for future use.

The old wrapping key is not deleted by this command.

 © Palo Alto Networks, Inc.

## Store Private Keys on an HSM

For added security, you can use an HSM to secure the private keys used in SSL/TLS decryption for:

- **SSL forward proxy**—The HSM can store the private key of the CA certificate that is used to sign certificates in SSL/TLS forward proxy operations. The firewall will then send the certificates that it generates during such operations to the HSM for signing before forwarding them to the client.
- **SSL inbound inspection**—The HSM can store the private keys for the internal servers for which you are performing SSL/TLS inbound inspection.

| Store Private Keys on an HSM | | |
|---|---|---|
| Step 1 | On the HSM, import or generate the private key used in your SSL forward proxy or SSL inbound inspection deployment. | For instructions on importing or generating a private key on the HSM, refer to your HSM documentation. |
| Step 2 | (Thales Nshield Connect only) Synchronize the key data from the HSM remote file system to the firewall. | 1. Access the firewall web interface and select **Device > Setup > HSM**.<br>2. Select **Synchronize with Remote Filesystem** in the Hardware Security Operations section. |
| Step 3 | Import the certificate that corresponds to the HSM-stored key onto the firewall. | 1. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>2. Enter the **Certificate Name**.<br>3. Enter the filename of the **Certificate File** you imported to the HSM.<br>4. Select a **File Format**.<br>5. Select the **Private Key resides on Hardware Security Module** check box.<br>6. Click **OK** and **Commit**. |
| Step 4 | (Forward trust certificates only) Enable the certificate for use in SSL/TLS Forward Proxy. | 1. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Open the certificate you imported in Step 3 for editing.<br>3. Select the **Forward Trust Certificate** check box.<br>4. Click **OK** and **Commit**. |
| Step 5 | Verify that you successfully imported the certificate onto the firewall. | 1. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Locate the certificate you imported in Step 3 and check the icon in the Key column:<br>• Lock icon—The private key for the certificate is on the HSM.<br>• Error icon—The private key is not on the HSM or the HSM is not properly authenticated or connected. |

© Palo Alto Networks, Inc.

## Manage the HSM Deployment

| Manage HSM | |
|---|---|
| • View the HSM configuration settings. | Select **Device > Setup > HSM**. |
| • Display detailed HSM information. | Select **Show Detailed Information** from the Hardware Security Operations section.<br>Information regarding the HSM servers, HSM HA status, and HSM hardware is displayed. |
| • Export Support file. | Select **Export Support File** from the Hardware Security Operations section.<br>A test file is created to help customer support when addressing a problem with an HSM configuration on the firewall. |
| • Reset HSM configuration. | Select **Reset HSM Configuration** from the Hardware Security Operations section.<br>Selecting this option removes all HSM connections. All authentication procedures must be repeated after using this option. |


© Palo Alto Networks, Inc.

# High Availability

High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point of failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up two firewalls in an HA pair provides redundancy and allows you to ensure business continuity.

The Palo Alto Networks firewalls support stateful active/passive or active/active high availability with session and configuration synchronization. Some models of the Palo Alto Networks firewall, such as the PA-200 only support HA lite without session synchronization capability, and the VM-Series firewall in AWS only supports active/passive HA. The following topics provide more information about high availability and how to configure it in your environment.

- HA Overview
- HA Concepts
- Set Up Active/Passive HA
- Reference: HA Synchronization
- HA Resources

# HA Overview

On Palo Alto Networks firewalls, you can set up two firewalls as an HA pair. HA allows you to minimize downtime by making sure that an alternate firewall is available in the event that the peer firewall fails. The firewalls in an HA pair use dedicated or in-band HA ports on the firewall to synchronize data—network, object, and policy configurations—and to maintain state information. Firewall-specific configuration such as management interface IP address or administrator profiles, HA specific configuration, log data, and the Application Command Center (ACC) information is not shared between peers. For a consolidated application and log view across the HA pair, you must use Panorama, the Palo Alto Networks centralized management system.

When a failure occurs on a firewall in an HA pair and the peer firewall takes over the task of securing traffic, the event is called a *failover*. The conditions that trigger a failover are:

- One or more of the monitored interfaces fail. (Link Monitoring)
- One or more of the destinations specified on the firewall cannot be reached. (Path Monitoring)
- The firewall does not respond to heartbeat polls. (Heartbeat Polling and Hello messages)

After you understand the HA Concepts, continue to Set Up Active/Passive HA.

 © Palo Alto Networks, Inc.

# HA Concepts

The following topics provide conceptual information about how HA works on a Palo Alto Networks firewall:

- HA Modes
- HA Links and Backup Links
- Device Priority and Preemption
- Failover Triggers
- HA Timers

## HA Modes

You can set up the firewalls for HA in two modes:

- **Active/Passive**— One firewall actively manages traffic while the other is synchronized and ready to transition to the active state, should a failure occur. In this configuration, both firewalls share the same configuration settings, and one actively manages traffic until a path, link, system, or network failure occurs. When the active firewall fails, the passive firewall transitions to the active state and takes over seamlessly and enforces the same policies to maintain network security. Active/passive HA is supported in the virtual wire, Layer 2 and Layer 3 deployments. For information on setting up your firewalls in an active/passive configuration, see Configure Active/Passive HA.

The PA-200 appliance only supports a lite version of active/passive HA. HA lite provides configuration synchronization and some runtime data synchronization such as IPSec security associations. It does not support any session synchronization, and therefore, HA Lite does not offer stateful failover.

- **Active/Active**— Both the firewalls in the pair are active and processing traffic, and work synchronously to handle session setup and session ownership. The active/active deployment is supported in virtual wire and Layer 3 deployments, and is only recommended for networks with asymmetric routing. For information on setting up the firewalls in an active/active configuration, refer to the Active/Active High Availability Tech Note.

## HA Links and Backup Links

The firewalls in an HA pair use HA links to synchronize data and maintain state information. Some models of the firewall have dedicated HA ports—Control link (HA1) and Data link (HA2), while others require you to use the in-band ports as HA links.

On firewalls with dedicated HA ports such as the PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7000 Series firewalls (see HA Ports on the PA-7000 Series Firewall), use the dedicated HA ports to manage communication and synchronization between the firewalls. For firewalls without dedicated HA ports such as the PA-200, PA-500, and PA-2000 Series firewalls, as a best practice use the management port for the HA1 link to allow for a direct connection between the management planes on the firewalls, and an in-band port for the HA2 link.

The HA1 and HA2 links provide synchronization for functions that reside on the management plane. Using the dedicated HA interfaces on the management plane is more efficient than using the in-band ports as this eliminates the need to pass the synchronization packets over the dataplane.

- **Control Link**: The HA1 link is used to exchange hellos, heartbeats, and HA state information, and management plane sync for routing, and User-ID information. The firewalls also use this link to synchronize configuration changes with its peer. The HA1 link is a Layer 3 link and requires an IP address.

  **Ports used for HA1**: TCP port 28769 and 28260 for clear text communication; port 28 for encrypted communication (SSH over TCP).

- **Data Link**: The HA2 link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between firewalls in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active or active-primary firewall to the passive or active-secondary firewall. The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.

  **Ports used for HA2:** The HA data link can be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets.

  Additionally, an HA3 link is used in Active/Active HA deployments. When there is an asymmetric route, the HA3 link is used for forwarding packets to the HA peer that owns the session. The HA3 link is a Layer 2 link and it does not support Layer 3 addressing or encryption.

- **Backup Links**: Provide redundancy for the HA1 and the HA2 links. In-band ports are used as backup links for both HA1 and HA2. Consider the following guidelines when configuring backup HA links:
  - The IP addresses of the primary and backup HA links must not overlap each other.
  - HA backup links must be on a different subnet from the primary HA links.
  - HA1-backup and HA2-backup ports must be configured on separate physical ports. The HA1-backup link uses port 28770 and 28260.

Palo Alto Networks recommends enabling heartbeat backup (uses port 28771 on the MGT interface) if you use an in-band port for the HA1 or the HA1 backup links.

- **Packet-Forwarding Link**: In addition to the HA1 and HA2 links, an active/active deployment also requires a dedicated HA3 link. The firewalls use this link for forwarding packets to the peer during session setup and asymmetric traffic flow. The HA3 link is a Layer 2 link that uses MAC-in-MAC encapsulation; it does not support Layer 3 addressing or encryption. You can configure aggregate interfaces on the PA-3000 Series, PA-4000 Series, PA-5000 Series and PA-7000 Series firewalls as an HA3 link. The aggregate interfaces can also provide redundancy for the HA3 link; you cannot configure backup links for the HA3 link.

## HA Ports on the PA-7000 Series Firewall

HA connectivity on the PA-7000 Series mandates the use of specific ports on the Switch Management Card (SMC) for certain functions; for other functions, you can use the ports on the Network Processing Card (NPC). The PA-7000 Series firewalls synchronize sessions across the NPCs one-for-one.

Use the following table for information on the SMC ports that are designed for HA connectivity:

 © Palo Alto Networks, Inc.

| HA Links and Backup Links | Ports on the SMC | Description |
|---|---|---|
| **Control Link** | HA1-A<br>Speed: Ethernet 10/100/1000 | Used for HA control and synchronization in both HA Modes. Connect this port directly from the HA1-A port on the first firewall to the HA1-A on the second firewall in the pair, or connect them together through a switch or router.<br>HA1 cannot be configured on NPC data ports or the MGT port. |
| **Control Link Backup** | HA1-B<br>Speed: Ethernet 10/100/1000 port | Used for HA control and synchronization as a backup for HA1-A in both HA Modes. Connect this port directly from the HA1-B port on the first firewall to the HA1-B on the second firewall in the pair, or connect them together through a switch or router.<br>HA1 Backup cannot be configured on NPC data ports or the MGT port. |
| **Data Link**<br><br>**Data Link Backup** | HSCI-A<br><br>HSCI-B | The High Speed Chassis Interconnect (HSCI) ports are Quad Port SFP (QSFP) interfaces which are used to connect two PA-7000 Series firewalls in an HA configuration. Each port is comprised of four 10 gigabit links internally for a combined speed of 40 gigabits.<br>The HSCI ports are not routable and must be connected directly to each other. The HSCI-A on the first chassis connects directly to HSCI-A on the second chassis and HSCI-B on the first chassis connects to HSCI-B on the second chassis. This will provide full 80 gigabit transfer rates. In software, both ports (HSCI-A and HSCI-B) are treated as one HA interface.<br>Palo Alto Networks recommends using the dedicated HSCI ports for the HA2 link; the HA3 link must use the HSCI port. If the firewalls are deployed in:<br>• an active/active configuration, the HA3 link must use the HSCI port. The HA2 link and HA2 backup links can use the HSCI port or data ports on the NPC.<br>• an active/passive configuration, you can configure a data port on the NPC for the HA2 link or the HA2 backup link, if needed. |

For an overview of the Modules and Interface cards on the PA-7000 Series firewall, refer to the PA-7000 Series Hardware Reference Guide.

## Device Priority and Preemption

The firewalls in an HA pair can be assigned a *device priority* value to indicate a preference for which firewall should assume the active or active-primary role. If you need to use a specific firewall in the HA pair for actively securing traffic, you must enable the preemptive behavior on both the firewalls and assign a device priority value for each firewall. The firewall with the lower numerical value, and therefore *higher priority*, is designated as active or active-primary. The other firewall is the active-secondary or passive firewall.

By default, preemption is disabled on the firewalls and must be enabled on both firewalls. When enabled, the preemptive behavior allows the firewall with the *higher priority* (lower numerical value) to resume as active or active-primary after it recovers from a failure. When preemption occurs, the event is logged in the system logs.

## Failover Triggers

When a failure occurs on one firewall and the peer takes over the task of securing traffic, the event is called a *failover*. A failover is triggered when a monitored metric on a firewall in the HA pair fails. The metrics that are monitored for detecting a firewall failure are:

- **Heartbeat Polling and Hello messages**

  The firewalls use hello message and heartbeats to verify that the peer firewall is responsive and operational. Hello messages are sent from one peer to the other at the configured *Hello Interval* to verify the state of the firewall. The heartbeat is an ICMP ping to the HA peer over the control link, and the peer responds to the ping to establish that the firewalls are connected and responsive. By default, the interval for the heartbeat is 1000 milliseconds. For details on the HA timers that trigger a failover, see HA Timers.

- **Link Monitoring**

  The physical interfaces to be monitored are grouped into a link group and their state (link up or link down) is monitored. A link group can contain one or more physical interfaces. A firewall failure is triggered when any or all of the interfaces in the group fail. The default behavior is failure of any one link in the link group will cause the firewall to change the HA state to non-functional to indicate a failure of a monitored object.

- **Path Monitoring**

  Monitors the full path through the network to mission-critical IP addresses. ICMP pings are used to verify reachability of the IP address. The default interval for pings is 200ms. An IP address is considered unreachable when 10 consecutive pings (the default value) fail, and a firewall failure is triggered when any or all of the IP addresses monitored become unreachable. The default behavior is any one of the IP addresses becoming unreachable will cause the firewall to change the HA state to non-functional to indicate a failure of a monitored object.

In addition to the failover triggers listed above, a failover also occurs when the administrator places the firewall is a suspended state or if preemption occurs.

On the PA-3000 Series, PA-5000 Series, and PA-7000 Series firewalls, a failover can occur when an internal health check fails. This health check is not configurable and is enabled to verify the operational status for all the components within the firewall.

 © Palo Alto Networks, Inc.

## HA Timers

High availability (HA) timers are used to detect a firewall failure and trigger a failover. To reduce the complexity in configuring HA timers, you can select from three profiles: **Recommended**, **Aggressive** and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.

Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

The following table describes each timer included in the profiles and the current preset values across the different hardware models; these values are for current reference only and can change in a subsequent release.

**Recommended/Aggressive HA Timer Values by Platform**

| Timers | Description | PA-7000 Series<br>PA-5000 Series<br>PA-4000 Series<br>PA-3000 Series<br>VM-Series | PA-2000 Series<br>PA-500 Series<br>PA-200 Series | Panorama Virtual Appliance<br>Panorama M-Series |
|---|---|---|---|---|
| Monitor fail hold up time | Interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is recommended to avoid an HA failover due to the occasional flapping of neighboring devices. | 0/0 | 0/0 | 0/0 |
| Preemption hold time | Time that a passive or active-secondary firewall will wait before taking over as the active or active-primary firewall. | 1/1 | 1/1 | 1/1 |
| Heartbeat interval | Frequency at which the HA peers exchange heartbeat messages in the form of an ICMP (ping). | 1000/1000<br>2000/1000 (only for VM-Series in AWS) | 2000/1000 | 2000/1000 |

| Timers | Description | PA-7000 Series<br>PA-5000 Series<br>PA-4000 Series<br>PA-3000 Series<br>VM-Series | PA-2000 Series<br>PA-500 Series<br>PA-200 Series | Panorama Virtual Appliance<br>Panorama M-Series |
|---|---|---|---|---|
| Promotion hold time | Time that the passive firewall (in active/passive mode) or the active-secondary firewall (in active/active mode) will wait before taking over as the active or active-primary firewall after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made. | 2000/500 | 2000/500 | 2000/500 |
| Additional master hold up time | Time interval that is applied to the same event as Monitor Fail Hold Up Time (range 0-60000 ms, default 500 ms). The additional time interval is applied only to the active firewall in active/passive mode and to the active-primary firewall in active/active mode. This timer is recommended to avoid a failover when both firewalls experience the same link/path monitor failure simultaneously. | 500/500 | 500/500 | 7000/5000 |
| Hello interval | Interval in milliseconds between hello packets that are sent to verify that the HA functionality on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms. | 8000/8000 | 8000/8000 | 8000/8000 |

 © Palo Alto Networks, Inc.

| Timers | Description | PA-7000 Series<br>PA-5000 Series<br>PA-4000 Series<br>PA-3000 Series<br>VM-Series | PA-2000 Series<br>PA-500 Series<br>PA-200 Series | Panorama Virtual Appliance<br>Panorama M-Series |
|---|---|---|---|---|
| Maximum no. of flaps | A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. This value indicates the maximum number of flaps that are permitted before the firewall is determined to be suspended and the passive firewall takes over (range 0-16; default 3). | 3/3 | 3/3 | Not Applicable |

# Set Up Active/Passive HA


- Prerequisites for Active/Passive HA
- Configuration Guidelines for Active/Passive HA
- Configure Active/Passive HA
- Define HA Failover Conditions
- Verify Failover


 © Palo Alto Networks, Inc.

## Prerequisites for Active/Passive HA

To set up high availability on your Palo Alto Networks firewalls, you need a pair of firewalls that meet the following requirements:

- **The same model**—Both the firewalls in the pair must be of the same hardware model or virtual machine model.
- **The same PAN-OS version**—Both the firewalls should be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases. They must also both have the same multiple virtual systems capability (single or multi vsys).
- **The same type of interfaces**—Dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
  - Determine the IP address for the HA1 (control) connection between the HA peers. The HA1 IP address for both peers must be on the same subnet if they are directly connected or are connected to the same switch.

    For firewalls without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both firewalls. However, because the management ports will not be directly cabled between the peers, make sure that you have a route that connects these two interfaces across your network.
  - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use Layer 3 *only* if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
- **The same set of licenses**—Licenses are unique to each firewall and cannot be shared between the firewalls. Therefore, you must license both firewalls identically. If both firewalls do not have an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.

If you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration, it is recommended that you Reset the Firewall to Factory Default Settings on the new firewall. This will ensure that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary firewall to the newly introduced firewall with the clean config.

## Configuration Guidelines for Active/Passive HA

To set up an active (PeerA) passive (PeerB) pair in HA, you must configure some options identically on both firewalls and some independently (non-matching) on each firewall. These HA settings are not synchronized between the firewalls. For details on what is/is not synchronized, refer to HA Synchronization.

To proceed with the instructions on configuring the firewalls in HA, see Configure Active/Passive HA.

The following table lists the settings that you must configure identically on both firewalls:

| Identical Configuration Settings on PeerA and PeerB |
|---|
| • HA must be enabled on both firewalls.<br>• Both firewalls must have the same Group ID value. The Group ID value is used to create a virtual MAC address for all the configured interfaces. The format of the virtual MAC is 00-1B-17:00: xx: yy where<br>00-1B-17: vendor ID; 00: fixed; xx: HA group ID; yy: interface ID.<br>When a new active firewall takes over, Gratuitous ARP messages are sent from each of the connected interfaces of the new active member to inform the connected Layer 2 switches of the virtual MAC address' new location.<br>• If using in-band ports, the interfaces for the HA1 and HA2 links must be set to type HA.<br>• The HA mode must be set to Active Passive.<br>• If required, preemption must be enabled on both firewalls. The device priority value, however, must not be identical.<br>• If required, encryption on the HA1 link (for communication between the HA peers) must be configured on both firewalls.<br>• Based on the combination of HA1 and HA1 Backup ports you are using, use the following recommendations to decide whether you should enable heartbeat backup:<br>• HA1: Dedicated HA1 port<br>HA1 Backup: In-band port<br>Recommendation: Enable Heartbeat Backup<br>• HA1: Dedicated HA1 port<br>HA1 Backup: Management port<br>Recommendation: Do not enable Heartbeat Backup<br>• HA1: In-band port<br>HA1 Backup: In-band port<br>Recommendation: Enable Heartbeat Backup<br>• HA1: Management port<br>HA1 Backup: In-band port<br>Recommendation: Do not enable Heartbeat Backup |

The following table lists the settings that must be configured independently on each firewall:

 © Palo Alto Networks, Inc.

| Independent Configuration Settings | PeerA | PeerB |
|---|---|---|
| Control Link | IP address of the HA1 link configured on this firewall (PeerA). | IP address of the HA1 link configured on this firewall (PeerB). |
| | For firewalls without dedicated HA ports, use the management port IP address for the control link. | |
| Data Link<br>The data link information is synchronized between the firewalls after HA is enabled and the control link is established between the firewalls. | By default, the HA2 link uses Ethernet/Layer 2.<br>If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerA). | By default, the HA2 link uses Ethernet/Layer 2.<br>If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerB). |
| Device Priority (required, if preemption is enabled) | The firewall you plan to make active must have a lower numerical value than its peer. So, if Peer A is to function as the active firewall, keep the default value of 100 and increment the value on PeerB. | If PeerB is passive, set the device priority value to a number larger than that on PeerA. For example, set the value to 110. |
| Link Monitoring—Monitor one or more physical interfaces that handle vital traffic on this firewall and define the failure condition. | Select the physical interfaces on the firewall that you would like to monitor and define the failure condition (all or any) to trigger a failover. | Pick a similar set of physical interfaces that you would like to monitor on this firewall and define the failure condition (all or any) to trigger a failover. |
| Path Monitoring—Monitor one or more destination IP addresses that the firewall can use ICMP pings to ascertain responsiveness. | Define the failure condition (all or any), ping interval and the ping count. This is particularly useful for monitoring the availability of other interconnected networking devices. For example, monitor the availability of a router that connects to a server, connectivity to the server itself, or some other vital device that is in the flow of traffic.<br>Make sure that the node/device that you are monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a a path monitoring failure and trigger a failover. | Pick a similar set of devices or destination IP addresses that can be monitored for determining the failover trigger for PeerB. Define the failure condition (all or any), ping interval and the ping count. |

## Configure Active/Passive HA

The following procedure shows how to configure a pair of firewalls in an active/passive deployment as depicted in the following example topology.

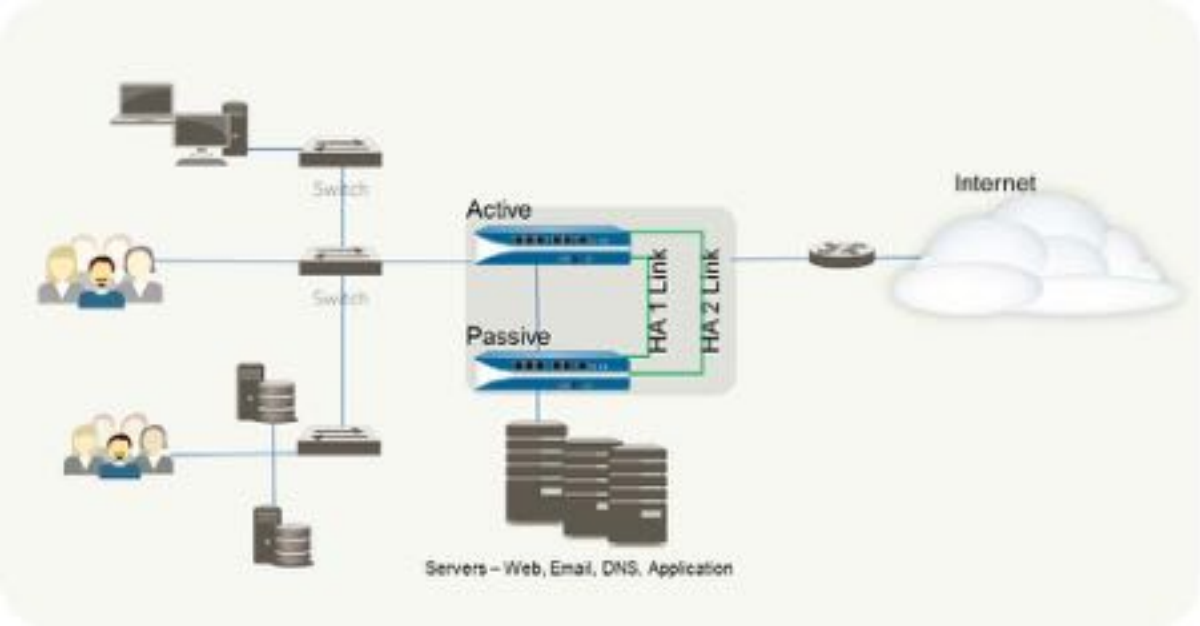


| Connect and Configure the Firewalls | | |
|---|---|---|
| Step 1 | Connect the HA ports to set up a physical connection between the firewalls. | • For firewalls with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on peers. Use a crossover cable if the peers are directly connected to each other.<br>• For firewalls without dedicated HA ports, select two data interfaces for the HA2 link and the backup HA1 link. Then, use an Ethernet cable to connect these in-band HA interfaces across both firewalls.<br>Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network. |
| Pick a firewall in the pair and complete the following steps: | | |
| Step 2 | Enable ping on the management port.<br>Enabling ping allows the management port to exchange heartbeat backup information. | 1. Select **Device > Setup > Management** and then click the Edit icon in the Management Interface Settings section of the screen.<br>2. Select **Ping** as a service that is permitted on the interface. |

 © Palo Alto Networks, Inc.

| Connect and Configure the Firewalls (Continued) | | |
|---|---|---|
| Step 3 | If the firewall does not have dedicated HA ports, set up the data ports to function as HA ports.<br><br>For firewalls with dedicated HA ports continue to the next step. | 1. Select **Network > Interfaces.**<br>2. Confirm that the link is up on the ports that you want to use.<br>3. Select the interface and set **Interface Type** to **HA**.<br> <br>4. Set the **Link Speed** and **Link Duplex** settings, as appropriate. |
| Step 4 | Set the HA mode and group ID. | 1. Select **Device > High Availability > General** and edit the Setup section.<br>2. Set a **Group ID** and optionally a **Description** for the pair. The Group ID uniquely identifies each HA pair on your network. If you have have multiple HA pairs that share the same broadcast domain you must set a unique Group ID for each pair.<br>3. Set the mode to **Active Passive**. |
| Step 5 | Set up the control link connection.<br><br>This example shows an in-band port that is set to interface type HA.<br><br>For firewalls that use the management port as the control link, the IP address information is automatically pre-populated. | 1. In **Device > High Availability > General**, edit the Control Link (HA1) section.<br>2. Select the **Port** that you have cabled for use as the HA1 link.<br>3. Set the **IPv4/IPv6 Address** and **Netmask**.<br>If the HA1 interfaces are on separate subnets, enter the IP address of the **Gateway**. Do not add a gateway address if the firewalls are directly connected<br> |
| Step 6 | (Optional) Enable encryption for the control link connection.<br><br>This is typically used to secure the link if the two firewalls are not directly connected, that is if the ports are connected to a switch or a router. | 1. Export the HA key from one firewall and import it into the peer firewall.<br>a. Select **Device > Certificate Management > Certificates**.<br>b. Select **Export HA key**. Save the HA key to a network location that the peer can access.<br>c. On the peer firewall, select **Device > Certificate Management > Certificates,** and select **Import HA key** to browse to the location that you saved the key and import it in to the peer.<br>2. Select **Device > High Availability > General**, edit the Control Link (HA1) section.<br>3. Select **Encryption Enabled**. |

| Connect and Configure the Firewalls (Continued) | |
|---|---|
| Step 7 Set up the backup control link connection. | 1. In **Device > High Availability > General**, edit the Control Link (HA1 Backup) section.<br>2. Select the HA1 backup interface and set the **IPv4/IPv6 Address** and **Netmask**.<br> |
| Step 8 Set up the data link connection (HA2) and the backup HA2 connection between the firewalls. | 1. In **Device > High Availability > General**, edit the Data Link (HA2) section.<br>2. Select the **Port** to use for the data link connection.<br>3. Select the **Transport** method. The default is **ethernet**, and will work when the HA pair is connected directly or through a switch. If you need to route the data link traffic through the network, select **IP** or **UDP** as the transport mode.<br>4. If you use IP or UDP as the transport method, enter the **IPv4/IPv6 Address** and **Netmask**..<br> 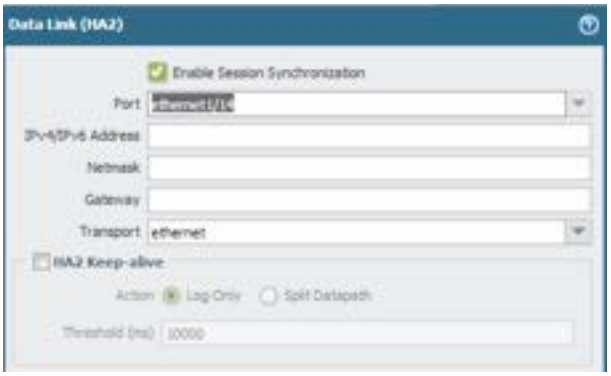 <br>5. Verify that **Enable Session Synchronization** is selected.<br>6. Select **HA2 Keep-alive** to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set (default is 10000 ms), the defined action will occur. For active/passive configuration, a critical system log message is generated when an HA2 keep-alive failure occurs.<br>You can configure the HA2 keep-alive option on both firewalls, or just one firewall in the HA pair. If the option is only enabled on one firewall, only that firewall will send the keep-alive messages. The other firewall will be notified if a failure occurs.<br>7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**. |

 © Palo Alto Networks, Inc.

| Connect and Configure the Firewalls (Continued) | | |
|---|---|---|
| Step 9 | Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port.<br><br>You do not need to enable heartbeat backup if you are using the management port for the control link. | 1. In **Device > High Availability > General**, edit the Election Settings.<br>2. Select **Heartbeat Backup**.<br>To allow the heartbeats to be transmitted between the firewalls, you must verify that the management port across both peers can route to each other.<br><br>Enabling heartbeat backup also allows you to prevent a split-brain situation. Split brain occurs when the HA1 link goes down causing the firewall to miss heartbeats, although the firewall is still functioning. In such a situation, each peer believes that the other is down and attempts to start services that are running, thereby causing a split brain. When the heartbeat backup link is enabled, split brain is prevented because redundant heartbeats and hello messages are transmitted over the management port. |
| Step 10 | Set the device priority and enable preemption.<br><br>This setting is only required if you wish to make sure that a specific firewall is the preferred active firewall. For information, see Device Priority and Preemption. | 1. In **Device > High Availability > General**, edit the Election Settings.<br>2. Set the numerical value in **Device Priority**. Make sure to set a lower numerical value on the firewall that you want to assign a higher priority to.<br><br>If both firewalls have the same device priority value, the firewall with the lowest MAC address on the HA1 control link will become the active firewall.<br>3. Select **Preemptive**.<br>You must enable preemptive on both the active firewall and the passive firewall. |
| Step 11 | (Optional) Modify the failover timers.<br><br>By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments. | 1. In **Device > High Availability > General**, edit the Election Settings.<br>2. Select the **Aggressive** profile for triggering failover faster; select **Advanced** to define custom values for triggering failover in your set up.<br><br>To view the preset value for an individual timer included in a profile, select **Advanced** and click **Load Recommended** or **Load Aggressive**. The preset values for your hardware model will be displayed on screen. |

| Connect and Configure the Firewalls (Continued) | |
|---|---|
| Step 12 (Optional, only configured on the passive firewall) Modify the link status of the HA ports on the passive firewall.<br><br>The passive link state is **shutdown**, by default. After you enable HA, the link state for the HA ports on the active firewall will be green and those on the passive firewall will be down and display as red. | Setting the link state to **Auto** allows for reducing the amount of time it takes for the passive firewall to take over when a failover occurs and it allows you to monitor the link state.<br><br>To enable the link status on the passive firewall to stay up and reflect the cabling status on the physical interface:<br>1. In **Device > High Availability > General**, edit the Active Passive Settings.<br>2. Set the **Passive Link State** to **Auto**.<br>The auto option decreases the amount of time it takes for the passive firewall to take over when a failover occurs.<br><br>Although the interface displays green (as cabled and up) it continues to discard all traffic until a failover is triggered.<br><br>When you modify the passive link state, make sure that the adjacent devices do not forward traffic to the passive firewall based only on the link status of the firewall. |
| Step 13 Enable HA. | 1. Select **Device > High Availability > General** and edit the Setup section.<br>2. Select **Enable HA**.<br>3. Select **Enable Config Sync**. This setting enables the synchronization of the configuration settings between the active and the passive firewall.<br>4. Enter the IP address assigned to the control link of the peer in **Peer HA1 IP Address**.<br> 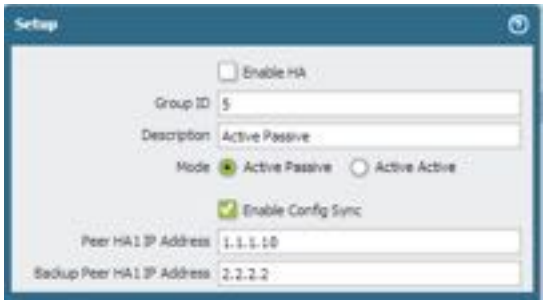 <br>For firewalls without dedicated HA ports, if the peer uses the management port for the HA1 link, enter the management port IP address of the peer.<br>5. Enter the **Backup HA1 IP Address**. |
| Step 14 Save your configuration changes. | Click **Commit**. |
| Step 15 Complete Step 2 through Step 14 on the other firewall in the HA pair. | |

 © Palo Alto Networks, Inc.

| Connect and Configure the Firewalls (Continued) | |
|---|---|
| Step 16 After you finish configuring both firewalls, verify that the firewalls are paired in active/passive HA. | 1. Access the **Dashboard** on both firewalls, and view the High Availability widget.<br>2. On the active firewall, click the **Sync to peer** link.<br>3. Confirm that the firewalls are paired and synced, as shown below: |
| On the passive firewall: the state of the local firewall should display **passive** and the Running Config should show as **synchronized**.<br> | On the active firewall: The state of the local firewall should display **active** and the Running Config should show as **synchronized**.<br> |

# Define HA Failover Conditions

| Configure the Failover Triggers | | |
|---|---|---|
| Step 1 | To configure link monitoring, define the interfaces that you would like to monitor. A change in the link state of these interface will trigger a failover. | 1. Select **Device > High Availability > Link and Path Monitoring** and **Add** a Link Group.<br>2. Name the **Link Group**, **Add** the interfaces to monitor, and select the **Failure Condition** for the group. The Link group you define is added to the **Link Group** section. |
| | | 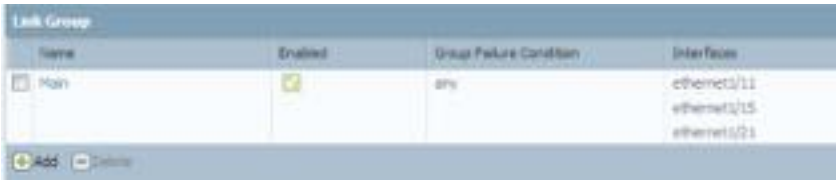 |
| Step 2 | (Optional) Modify the failure condition for the Link Groups that you configured (in the preceding step) on the firewall.<br>By default, the firewall will trigger a failover when any monitored link fails. | 1. Select the **Link Monitoring** section.<br>2. Set the **Failure Condition** to **All**.<br>The default setting is **Any**. |
| Step 3 | To configure path monitoring, define the destination IP addresses that the firewall should ping to verify network connectivity. | 1. In the **Path Group** section of the **Device > High Availability > Link and Path Monitoring** tab, pick the **Add option for your set up:** Virtual Wire, VLAN, or Virtual Router.<br>2. Select the appropriate item from the drop-down for the **Name** and **Add** the IP addresses (source and/or destination, as prompted) that you wish to monitor. Then select the **Failure Condition** for the group. The path group you define is added to the **Path Group** section.<br> |
| Step 4 | (Optional) Modify the failure condition for all Path Groups configured on the firewall.<br>By default, the firewall will trigger a failover when any monitored path fails. | Set the **Failure Condition** to **All.**<br>The default setting is **Any**. |
| Step 5 | Save your changes. | Click **Commit**. |

 © Palo Alto Networks, Inc.

If you are using SNMPv3 to monitor the firewalls, note that the SNMPv3 Engine ID is unique to each firewall; the EngineID is not synchronized between the HA pair and, therefore, allows you to independently monitor each firewall in the HA pair. For information on setting up SNMP, see Forward Traps to an SNMP Manager.

Because the EngineID is generated using the firewall serial number, on the VM-Series firewall you must apply a valid license in order to obtain a unique EngineID for each firewall.

## Verify Failover

To test that your HA configuration works properly, trigger a manual failover and verify that the firewalls transition states successfully.

| Verify Failover | | |
|---|---|---|
| Step 1 | Suspend the active firewall. | Select **Device > High Availability > Operational Commands** and click the **Suspend local device** link.<br> 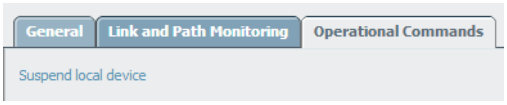 |
| Step 2 | Verify that the passive firewall has taken over as active. | On the **Dashboard**, verify that the state of the passive firewall changes to **active** in the High Availability widget.<br> |
| Step 3 | Restore the suspended firewall to a functional state. Wait for a couple minutes, and then verify that preemption has occurred, if preemptive is enabled. | 1. On the firewall you previously suspended, select **Device > High Availability > Operational Commands** and click the **Make local device functional** link.<br> 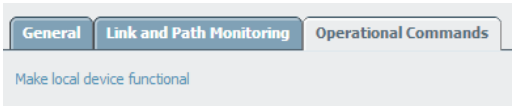 <br>2. In the High Availability widget on the **Dashboard,** confirm that the firewall has taken over as the active firewall and that the peer is now in a passive state.<br> |

 © Palo Alto Networks, Inc.

# Reference: HA Synchronization

If you have enabled configuration synchronization on both peers in an HA pair, most of the configuration settings you configure on one peer will automatically sync to the other peer upon commit. To avoid configuration conflicts, always make configuration changes on the active (active/passive) or active-primary (active/active) peer and wait for the changes to sync to the peer before making any additional configuration changes.

The following topics identify what portions of a firewall configuration must be configured on each device independently (rather than synchronized from the HA peer).


- What Settings Don't Sync in Active/Passive HA?
- What Settings Don't Sync in Active/Active HA?
- Synchronization of System Runtime Information


## What Settings Don't Sync in Active/Passive HA?

You must configure the following settings on each firewall in an HA pair in an active/passive deployment. These settings do not sync from one peer to another:

| Configuration Item | What Doesn't Sync in Active/Passive? |
|---|---|
| Management Interface Settings | All management configuration settings must be configured individually on each device, including:<br>• **Device > Setup > Management > General Settings**—Hostname, Domain, Login Banner, Time Zone, Locale, Date, Time, Latitude, Longitude<br>• **Device > Setup > Management > Management Interface Settings**—IP Address, Netmask, Default Gateway, IPv6 Address/Prefix Length, Default IPv6 Gateway, Speed, MTU, and Services (HTTP, HTTP OCSP, HTTPS, Telnet, SSH, Ping, SNMP, User-ID, User-ID Syslog Listener-SSL, User-ID Syslog Listener-UDP) |
| Multi-vsys Capability | To enable multi-vsys you must activate the Virtual Systems license (required to enable support for multiple virtual systems on PA-2000 Series and PA-3000 Series firewalls or to increase the number of virtual systems beyond the base number provided by default on PA-4000 Series, PA-5000 Series, and PA-7000 Series firewalls) on each firewall in the pair.<br>In addition, you must also enable **Multi Virtual System Capability** on each firewall (**Device > Setup > Management > General Settings**). |
| Administrator Authentication Settings | You must define the authentication profile and certificate profile for administrative access to the firewall locally on each firewall (**Device > Setup > Management > Authentication**). |
| Panorama Settings | Set the following Panorama settings on each firewall (**Device > Setup > Management > Panorama Settings**).<br>• **Panorama Servers**<br>• **Disable Panorama Policy and Objects** and **Disable Device and Network Template** |

| Configuration Item | What Doesn't Sync in Active/Passive? |
|---|---|
| SNMP | • **Device > Setup > Operations > SNMP Setup** |
| Statistics Collection | • **Device > Setup > Operations > Statistics Service Setup** |
| Services | • **Device > Setup > Services** |
| Global Service Routes | • **Device > Setup > Services > Service Route Configuration** |
| Data Protection | • **Device > Setup > Content-ID > Manage Data Protection** |
| Jumbo Frames | • **Device > Setup > Session > Session Settings > Enable Jumbo Frame** |
| Forward Proxy Server Certificate Settings | • **Device > Setup > Session > Decryption Settings > SSL Forward Proxy Settings** |
| Master Key Secured by HSM | • **Device > Setup > HSM > Hardware Security Module Provider > Master Key Secured by HSM** |
| Log Export Settings | • **Device > Scheduled Log Export** |
| Software Updates | With software updates, you can either download and install them separately on each device, or download them on one peer and sync the update to the other peer. You must install the update on each peer.<br>• **Device > Software** |
| GlobalProtect Agent Package | With GlobalProtect client updates, you can either download and install them separately on each device, or download them to one peer and sync the update to the other peer. You must activate separately on each peer.<br>• **Device > GlobalProtect Client** |
| Content Updates | With content updates, you can either download and install them separately on each device, or download them on one peer and sync the update to the other peer. You must install the update on each peer.<br>• **Device > Dynamic Updates** |
| Licenses/Subscriptions | • **Device > Licenses** |
| Support Subscription | • **Device > Support** |
| Master Key | The master key must be identical on each firewall in the HA pair, but you must manually enter it on each device (**Device > Master Key and Diagnostics**).<br>Before changing the master key, you must disable config sync on both peers (**Device > High Availability > General > Setup** and clear the **Enable Config Sync** check box) and then re-enable it after you change the keys. |
| Reports, logs, and Dashboard Settings | Log data, reports, and Dashboard data and settings (column display, widgets) are not synced between peers. Report configuration settings, however, are synced. |

 © Palo Alto Networks, Inc.

## What Settings Don't Sync in Active/Active HA?

You must configure the following settings on each firewall in an HA pair in an active/active deployment. These settings do not sync from one peer to another:

| Configuration Item | What Doesn't Sync? |
|---|---|
| Management Interface Settings | All management configuration settings must be configured individually on each device, including:<br>• **Device > Setup > Management > General Settings**—Hostname, Domain, Login Banner, Time Zone, Locale, Date, Time, Latitude, Longitude<br>• **Device > Setup > Management > Management Interface Settings**—IP Address, Netmask, Default Gateway, IPv6 Address/Prefix Length, Default IPv6 Gateway, Speed, MTU, and Services (HTTP, HTTP OCSP, HTTPS, Telnet, SSH, Ping, SNMP, User-ID, User-ID Syslog Listener-SSL, User-ID Syslog Listener-UDP) |
| Multi-vsys Capability | To enable multi-vsys you must activate the Virtual Systems license (required to enable support for multiple virtual systems on PA-2000 Series and PA-3000 Series firewalls or to increase the number of virtual systems beyond the base number provided by default on PA-4000 Series, PA-5000 Series, and PA-7000 Series firewalls) on each firewall in the pair.<br>In addition, you must also enable **Multi Virtual System Capability** on each firewall (**Device > Setup > Management > General Settings**). |
| Administrator Authentication Settings | You must define the authentication profile and certificate profile for administrative access to the firewall locally on each firewall (**Device > Setup > Management > Authentication**). |
| Panorama Settings | Set the following Panorama settings on each firewall (**Device > Setup > Management > Panorama Settings**).<br>• **Panorama Servers**<br>• **Disable Panorama Policy and Objects** and **Disable Device and Network Template** |
| SNMP | • **Device > Setup > Operations > SNMP Setup** |
| Statistics Collection | • **Device > Setup > Operations > Statistics Service Setup** |
| Services | • **Device > Setup > Services** |
| Global Service Routes | • **Device > Setup > Services > Service Route Configuration** |
| Data Protection | • **Device > Setup > Content-ID > Manage Data Protection** |
| Jumbo Frames | • **Device > Setup > Session > Session Settings > Enable Jumbo Frame** |
| Forward Proxy Server Certificate Settings | • **Device > Setup > Session > Decryption Settings > SSL Forward Proxy Settings** |
| HSM Configuration | • **Device > Setup > HSM** |
| Log Export Settings | • **Device > Scheduled Log Export** |

| Configuration Item | What Doesn't Sync? |
|---|---|
| Software Updates | With software updates, you can either download and install them separately on each device, or download them on one peer and sync the update to the other peer. You must install the update on each peer.<br>• **Device > Software** |
| GlobalProtect Agent Package | With GlobalProtect client updates, you can either download and install them separately on each device, or download them to one peer and sync the update to the other peer. You must activate separately on each peer.<br>• **Device > GlobalProtect Client** |
| Content Updates | With content updates, you can either download and install them separately on each device, or download them on one peer and sync the update to the other peer. You must install the update on each peer.<br>• **Device > Dynamic Updates** |
| Licenses/Subscriptions | • **Device > Licenses** |
| Support Subscription | • **Device > Support** |
| Ethernet Interface IP Addresses | All Ethernet interface configuration settings sync except for the IP address (**Network > Interface > Ethernet**). |
| Loopback Interface IP Addresses | All Loopback interface configuration settings sync except for the IP address (**Network > Interface > Loopback**). |
| Tunnel Interface IP Addresses | All Tunnel interface configuration settings sync except for the IP address (**Network > Interface > Tunnel**). |
| LACP System Priority | Each peer must have a unique LACP System ID in an active/active deployment (**Network > Interface > Ethernet > Add Aggregate Group > System Priority**). |
| VLAN Interface IP Address | All VLAN interface configuration settings sync except for the IP address (**Network > Interface > VLAN**). |
| Virtual Routers | Virtual router configuration synchronizes only if you have enabled VR Sync (**Device > High Availability > Active/Active Config > Packet Forwarding**). Whether or not to do this depends on your network design, including whether you have asymmetric routing. |
| IPSec Tunnels | IPSec tunnel configuration synchronization is dependent on whether you have configured the Virtual Addresses to use Floating IP addresses (**Device > High Availability > Active/Active Config > Virtual Address**). If you have configured a floating IP address, these settings sync automatically. Otherwise, you must configure these settings independently on each peer. |
| GlobalProtect Portal Configuration | GlobalProtect portal configuration synchronization is dependent on whether you have configured the Virtual Addresses to use Floating IP addresses (**Network > GlobalProtect > Portals**). If you have configured a floating IP address, the GlobalProtect portal configuration settings sync automatically. Otherwise, you must configure the portal settings independently on each peer. |

 © Palo Alto Networks, Inc.

| Configuration Item | What Doesn't Sync? |
|---|---|
| GlobalProtect Gateway Configuration | GlobalProtect gateway configuration synchronization is dependent on whether you have configured the Virtual Addresses to use Floating IP addresses (**Network > GlobalProtect > Gateways**). If you have configured a floating IP address, the GlobalProtect gateway configuration settings sync automatically. Otherwise, you must configure the gateway settings independently on each peer. |
| QoS | QoS configuration synchronizes only if you have enabled **QoS Sync** (**Device > High Availability > Active/Active Config > Packet Forwarding**). You might choose not to sync QoS setting if, for example, you have different bandwidth on each link or different latency through your service providers. |
| LLDP | No LLDP state or individual firewall data is synchronized in an active/active configuration (**Network > LLDP**). |
| IKE Gateways | IKE gateway configuration synchronization is dependent on whether you have configured the Virtual Addresses to use floating IP addresses (**Network > IKE Gateways**). If you have configured a floating IP address, the IKE gateway configuration settings sync automatically. Otherwise, you must configure the IKE gateway settings independently on each peer. |
| Master Key | The master key must be identical on each firewall in the HA pair, but you must manually enter it on each device (**Device > Master Key and Diagnostics**).<br>Before changing the master key, you must disable config sync on both peers (**Device > High Availability > General > Setup** and clear the **Enable Config Sync** check box) and then re-enable it after you change the keys. |
| Reports, logs, and Dashboard Settings | Log data, reports, and dashboard data and settings (column display, widgets) are not synced between peers. Report configuration settings, however, are synced. |

## Synchronization of System Runtime Information

| Runtime Information | Config Synced? | | HA Link | Details |
|---|---|---|---|---|
| | A/P | A/A | | |
| **Management Plane** | | | | |
| User to Group Mappings | Yes | Yes | HA1 | |
| DHCP Lease (as server) | Yes | Yes | HA1 | |
| DNS Cache | Yes | Yes | HA1 | |
| FQDN Refresh | No | No | HA1 | |
| IKE Keys (phase 2) | Yes | Yes | HA1 | |
| BrightCloud URL Database | No | No | N/A | |

| Runtime Information | Config Synced? | | HA Link | Details |
|---|---|---|---|---|
| | A/P | A/A | | |
| BrightCloud URL Cache | No | No | N/A | This feature is disabled by default and must be enabled separately on each HA peer. |
| BrightCloud Bloom Filter | No | No | N/A | This feature is disabled by default and must be enabled separately on each HA peer. |
| PAN-DB URL Cache | Yes | No | HA1 | This is synchronized upon database backup to disk (every eight hours, when URL database version updates), or when the firewall reboots. |
| Content (manual sync) | Yes | Yes | HA1 | |
| PPPoE, PPPoE Lease | Yes | Yes | HA1 | |
| DHCP Client Settings and Lease | Yes | Yes | HA1 | |
| SSL VPN Logged in User List | Yes | Yes | HA1 | |
| Forward Information Base (FIB) | Yes | Yes | HA1 | |
| **Dataplane** | | | | |
| Session Table | Yes | Yes | HA2 | • Active/passive peers do not sync ICMP or host session information.<br>• Active/active peers do not sync host session or multicast session information. |
| ARP Table | Yes | No | HA2 | |
| Neighbor Discovery (ND) Table | Yes | No | HA2 | |
| MAC Table | Yes | No | HA2 | |
| IPSec Sequence Number (anti-replay) | Yes | Yes | HA2 | |
| DoS Protection | Yes | Yes | HA2 | |
| User to IP Address Mappings | Yes | Yes | HA2 | |
| Virtual MAC | Yes | Yes | HA2 | |

 © Palo Alto Networks, Inc.

# HA Resources

For more information on HA, refer to the following sources:

- Active/Active HA
- High Availability Failover Optimization
- Upgrading an HA pair
- Examples: Deploying HA

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

# Monitoring

To forestall potential issues, and accelerate incidence response when needed, the firewall provides intelligence on traffic and user patterns and customizable and informative reports. The dashboard, Application Command Center (ACC), reports, and logs on the firewall allow you to monitor activity on your network. You can monitor the logs and filter the information to generate reports with predefined or customized views. You can, for example, use the predefined templates to generate reports on user activities, or analyze the reports and logs to interpret unusual behavior on your network and generate a custom report on the traffic pattern. For a visually engaging presentation of network activity, the dashboard and the ACC include widgets, charts, and tables that you can interact with to find information that you care about. In addition, you can configure the firewall to forward monitored information as email notifications, syslog messages, SNMP traps, and NetFlow records to external services.


- Use the Dashboard
- Use the Application Command Center
- App Scope
- Use the Automated Correlation Engine
- Take Packet Captures
- Monitor Applications and Threats
- Monitor and Manage Logs
- Manage Reporting
- Use External Services for Monitoring
- Configure Log Forwarding
- Configure Email Alerts
- Use Syslog for Monitoring
- SNMP Monitoring and Traps
- NetFlow Monitoring


© Palo Alto Networks, Inc.

# Use the Dashboard

The **Dashboard** tab widgets show general device information, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available widgets are displayed by default, but each administrator can remove and add individual widgets, as needed. Click the refresh icon to update the dashboard or an individual widget. To change the automatic refresh interval, select an interval from the drop-down (**1 min**, **2 mins**, **5 mins**, or **Manual**). To add a widget to the dashboard, click the widget drop-down, select a category and then the widget name. To delete a widget, click in the title bar. The following table describes the dashboard widgets.

| Dashboard Charts | Descriptions |
|---|---|
| Top Applications | Displays the applications with the most sessions. The block size indicates the relative number of sessions (mouse-over the block to view the number), and the color indicates the security risk—from green (lowest) to red (highest). Click an application to view its application profile. |
| Top High Risk Applications | Similar to Top Applications, except that it displays the highest-risk applications with the most sessions. |
| General Information | Displays the device name, model, PAN-OS software version, the application, threat, and URL filtering definition versions, the current date and time, and the length of time since the last restart. |
| Interface Status | Indicates whether each interface is up (green), down (red), or in an unknown state (gray). |
| Threat Logs | Displays the threat ID, application, and date and time for the last 10 entries in the Threat log. The threat ID is a malware description or URL that violates the URL filtering profile. |
| Config Logs | Displays the administrator username, client (Web or CLI), and date and time for the last 10 entries in the Configuration log. |
| Data Filtering Logs | Displays the description and date and time for the last 60 minutes in the Data Filtering log. |
| URL Filtering Logs | Displays the description and date and time for the last 60 minutes in the URL Filtering log. |
| System Logs | Displays the description and date and time for the last 10 entries in the System log.<br>A `Config installed` entry indicates configuration changes were committed successfully. |
| System Resources | Displays the Management CPU usage, Data Plane usage, and the Session Count, which displays the number of sessions established through the firewall. |
| Logged In Admins | Displays the source IP address, session type (Web or CLI), and session start time for each administrator who is currently logged in. |
| ACC Risk Factor | Displays the average risk factor (1 to 5) for the network traffic processed over the past week. Higher values indicate higher risk. |
| High Availability | If high availability (HA) is enabled, indicates the HA status of the local and peer device—green (active), yellow (passive), or black (other). For more information about HA, see High Availability. |
| Locks | Shows configuration locks taken by administrators. |

 © Palo Alto Networks, Inc.

# Use the Application Command Center

The Application Command Center (ACC) is an interactive, graphical summary of the applications, users, URLs, threats, and content traversing your network.The ACC uses the firewall logs to provide visibility into traffic patterns and actionable information on threats. The ACC layout includes a tabbed view of network activity, threat activity, and blocked activity and each tab includes pertinent widgets for better visualization of network traffic. The graphical representation allows you to interact with the data and visualize the relationships between events on the network, so that you can uncover anomalies or find ways to enhance your network security rules. For a personalized view of your network, you can also add a custom tab and include widgets that allow you to drill down into the information that is most important to you.


- ACC—First Look
- ACC Tabs
- ACC Widgets (Widget Descriptions)
- ACC Filters
- Interact with the ACC
- Use Case: ACC—Path of Information Discovery

# ACC—First Look

Take a quick tour of the ACC.

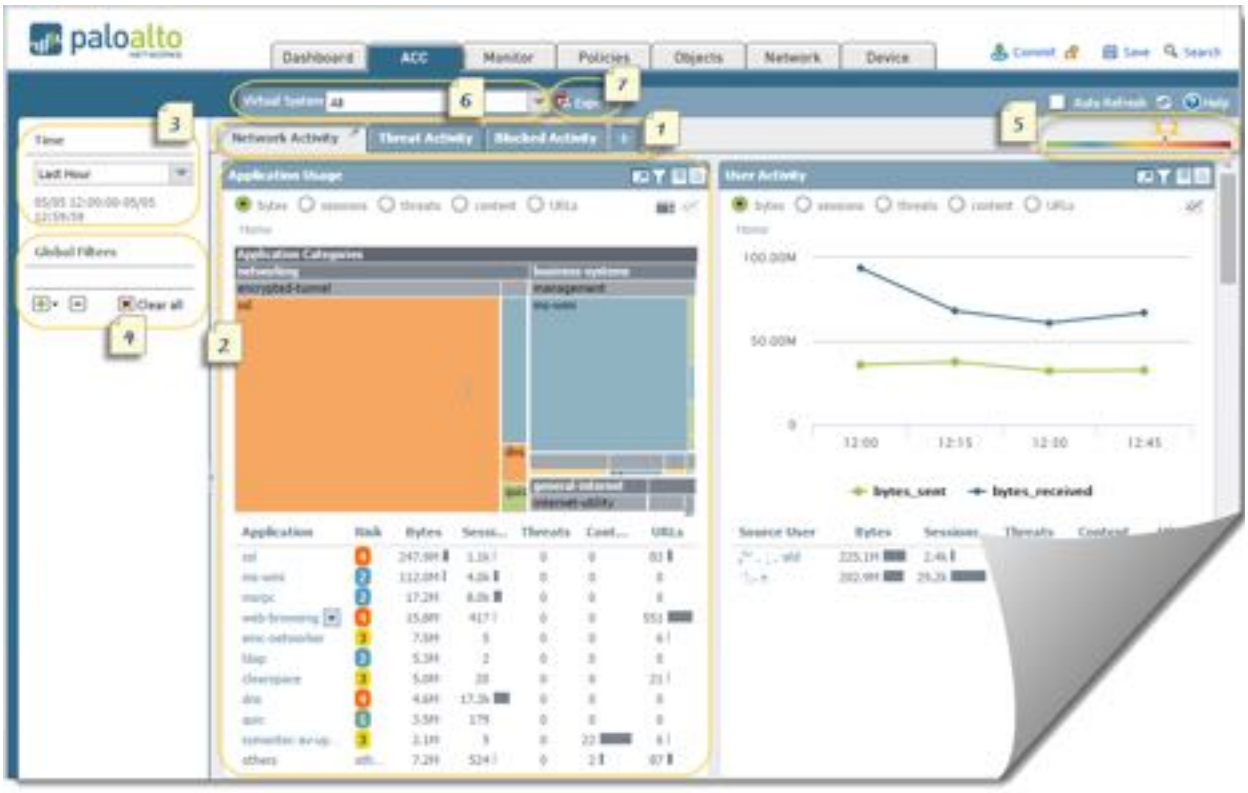


| ACC—First Look | | |
|---|---|---|
| 1 | **Tabs** | The ACC includes three predefined tabs that provide visibility into network traffic, threat activity, and blocked activity. For information on each tab, see ACC Tabs. |
| 2 | **Widgets** | Each tab includes a default set of widgets that best represent the events/trends associated with the tab. The widgets allow you to survey the data using the following filters:<br>• bytes (in and out)<br>• sessions<br>• content (files and data)<br>• URL categories<br>• threats (and count)<br>For information on each widget, see ACC Widgets. |


© Palo Alto Networks, Inc.

| ACC—First Look (Continued) | | |
|---|---|---|
| 3 | **Time** | The charts or graphs in each widget provide a summary and historic view. You can choose a custom range or use the predefined time periods that range from the last 15 minutes up to the last 30 days or last 30 calendar days. The selected time period applies across all tabs in the ACC.<br>The time period used to render data, by default, is the **Last Hour** updated in 15 minute intervals. The date and time interval are displayed onscreen, for example at 11:40, the time range is 01/12 10:30:00-01/12 11:29:59. |
| 4 | **Global Filters** | The Global Filters allow you to set the filter across all widgets and all tabs. The charts/graphs apply the selected filters before rendering the data. For information on using the filters, see ACC Filters. |
| 5 | **Risk Factor** | The risk factor (1=lowest to 5=highest) indicates the relative risk based on the applications used on your network. The risk factor uses a variety of factors to assess the associated risk levels, such as whether the application can share files, is it prone to misuse or does it try to evade firewalls, it also factors in the threat activity and malware as seen through the number of blocked threats, compromised hosts or traffic to malware hosts/domains. |
| 6 | **Source** | The data segment used for the display. The options vary on the firewall and on Panorama.<br>On the firewall, if enabled for multiple virtual systems, you can use the **Virtual System** drop-down to change the ACC display to include all virtual systems or just a selected virtual system.<br>On Panorama, you can select the **Device Group** drop-down to change the ACC display to include all device groups or just a selected device group.<br>Additionally, on Panorama, you can change the **Data Source** as **Panorama** data or **Remote Device Data**. **Remote Device Data** is only available when all the managed firewalls are on PAN-OS 7.0.0 or later. When you filter the display for a specific device group, **Panorama** data is used as the data source**.** |
| 7 | **Export** | You can export the widgets displayed in the currently selected tab as a PDF. The PDF is downloaded and saved to the downloads folder associated with your web browser, on your computer. |

## ACC Tabs

The ACC includes the following predefined tabs for viewing network activity, threat activity, and blocked activity.

| Tab | Description |
|---|---|
| **Network Activity** | Displays an overview of traffic and user activity on your network including:<br>• Top applications in use<br>• Top users who generate traffic (with a drill down into the bytes, content, threats or URLs accessed by the user)<br>• Most used security rules against which traffic matches occur<br>In addition, you can also view network activity by source or destination zone, region, or IP address, ingress or egress interfaces, and GlobalProtect host information such as the operating systems of the devices most commonly used on the network. |
| **Threat Activity** | Displays an overview of the threats on the network, focusing on the top threats: vulnerabilities, spyware, viruses, hosts visiting malicious domains or URLs, top WildFire submissions by file type and application, and applications that use non-standard ports. The Compromised Hosts widget in this tab (the widget is supported on some platforms only), supplements detection with better visualization techniques; it uses the information from the correlated events tab (**Automated Correlation Engine > Correlated Events**) to present an aggregated view of compromised hosts on your network by source users/IP addresses and sorted by severity. |
| **Blocked Activity** | Focuses on traffic that was prevented from coming into the network. The widgets in this tab allow you to view activity denied by application name, username, threat name, blocked content—files and data that were blocked by a file blocking profile. It also lists the top security rules that were matched on to block threats, content, and URLs. |

You can also Interact with the ACC to create customized tabs with custom layout and widgets that meet your network monitoring needs.

 © Palo Alto Networks, Inc.

## ACC Widgets

The widgets on each tab are interactive; you can set the ACC Filters and drill down into the details for each table or graph, or customize the widgets included in the tab to focus on the information you need. For details on what each widget displays, see Widget Descriptions.

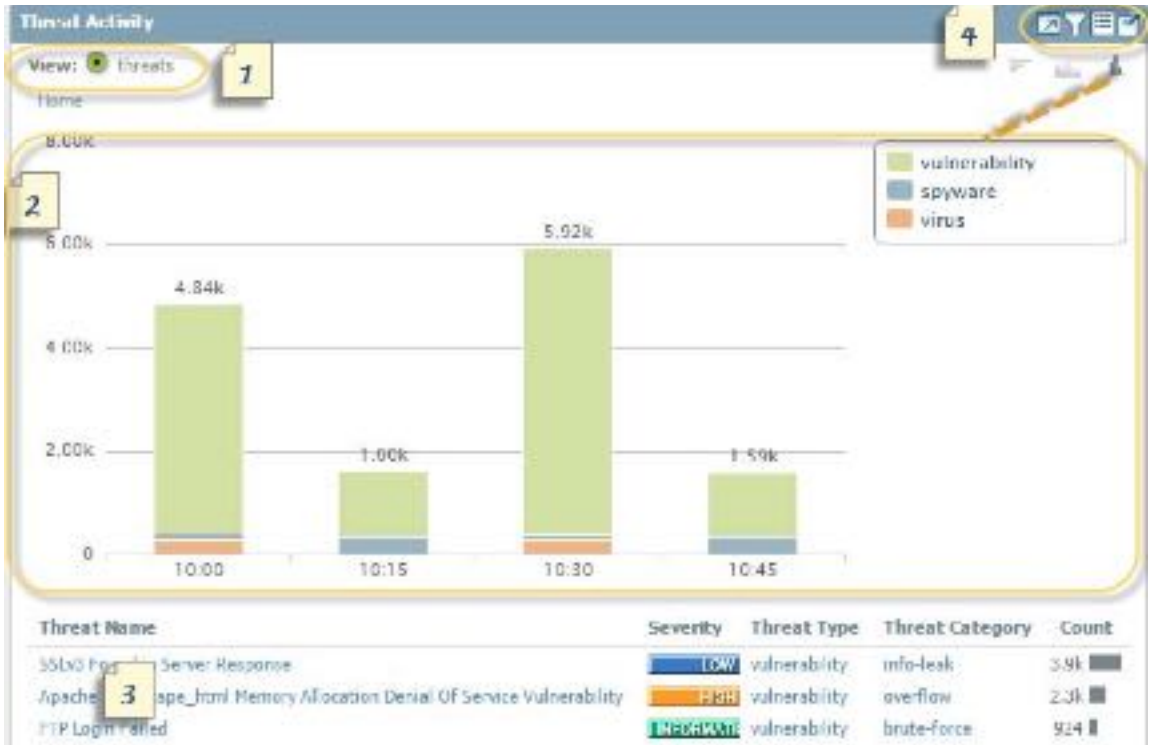


| Widgets | | |
|---|---|---|
| 1 | **View** | You can sort the data by bytes, sessions, threats, count, content, URLs, malicious, benign, files, data, profiles, objects. The available options vary by widget. |
| 2 | **Graph** | The graphical display options are treemap, line graph, horizontal bar graph, stacked area graph, stacked bar graph, and map. The available options vary by widget; the interaction experience also varies with each graph type. For example, the widget for Applications using Non-Standard Ports allows you to choose between a treemap and a line graph.<br>To drill down into the display, click into the graph. The area you click into becomes a filter and allows you to zoom into the selection and view more granular information on the selection. |

| Widgets | | |
|---|---|---|
| 3 | **Table** | The detailed view of the data used to render the graph is provided in a table below the graph. You can interact with the table in several ways:<br>• Click and set a local filter for an attribute in the table. The graph is updated and the table is sorted using the local filter. The information displayed in the graph and the table are always synchronized.<br>• Hover over the attribute in the table and use the options available in the drop-down.<br> |
| 4 | **Actions** | **Maximize view**— Allows you enlarge the widget and view the table in a larger screen space and with more viewable information.<br>**Set up local filters**—Allows you to add ACC Filters to refine the display within the widget. Use these filters to customize the widgets; these customizations are retained between logins.<br>**Jump to logs**—Allows you to directly navigate to the logs (**Monitor > Logs >** *Log type* tab). The logs are filtered using the time period for which the graph is rendered. If you have set local and global filters, the log query concatenates the time period and the filters and only displays logs that match the combined filter set.<br>**Export**—Allows you to export the graph as a PDF. The PDF is downloaded and saved on your computer. It is saved in the Downloads folder associated with your web browser. |

 © Palo Alto Networks, Inc.

## Widget Descriptions

Each tab on the ACC includes a different set of widgets.

| Widget | Description |
|---|---|
| **Network Activity**—Displays an overview of traffic and user activity on your network. | |
| **Application Usage** | The table displays the top ten applications used on your network, all the remaining applications used on the network are aggregated and displayed as other. The graph displays all applications by application category, sub category, and application. Use this widget to scan for applications being used on the network, it informs you about the predominant applications using bandwidth, session count, file transfers, triggering the most threats, and accessing URLs.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: treemap, area, column, line (the charts vary by the sort by attribute selected) |
| **User Activity** | Displays the top ten most active users on the network who have generated the largest volume of traffic and consumed network resources to obtain content. Use this widget to monitor top users on usage sorted on bytes, sessions, threats, content (files and patterns), and URLs visited.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: area, column, line (the charts vary by the sort by attribute selected) |
| **Source IP Activity** | Displays the top ten IP addresses or hostnames of the devices that have initiated activity on the network. All other devices are aggregated and displayed as other.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: area, column, line (the charts vary by the sort by attribute selected) |
| **Destination IP Activity** | Displays the IP addresses or hostnames of the top ten destinations that were accessed by users on the network.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: area, column, line (the charts vary by the sort by attribute selected) |
| **Source Regions** | Displays the top ten regions (built-in or custom defined regions) around the world from where users initiated activity on your network.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: map, bar |
| **Destination Regions** | Displays the top ten destination regions (built-in or custom defined regions) on the world map from where content is being accessed by users on the network.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: map, bar |

| Widget | Description |
|---|---|
| **GlobalProtect Host Information** | Displays information on the state of the hosts on which the GlobalProtect agent is running; the host system is a GlobalProtect client. This information is sourced from entries in the HIP match log that are generated when the data submitted by the GlobalProtect agent matches a HIP object or a HIP profile you have defined on the firewall. If you do not have HIP Match logs, this widget is blank. To learn how to create HIP objects and HIP profiles and use them as policy match criteria, see Configure HIP-Based Policy Enforcement.<br>Sort attributes: profiles, objects, operating systems<br>Charts available: bar |
| **Rule Usage** | Displays the top ten rules that have allowed the most traffic on the network. Use this widget to view the most commonly used rules, monitor the usage patterns, and to assess whether the rules are effective in securing your network.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: line |
| **Ingress Interfaces** | Displays the firewall interfaces that are most used for allowing traffic into the network.<br>Sort attributes: bytes, bytes sent, bytes received<br>Charts available: line |
| **Egress Interfaces** | Displays the firewall interfaces that are most used by traffic exiting the network.<br>Sort attributes: bytes, bytes sent, bytes received<br>Charts available: line |
| **Source Zones** | Displays the zones that are most used for allowing traffic into the network.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: line |
| **Destination Zones** | Displays the zones that are most used by traffic going outside the network.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: line |

**Threat Activity**—Displays an overview of the threats on the network

| | |
|---|---|
| **Compromised Hosts** | Displays the hosts that are likely compromised on your network. This widget summarizes the events from the correlation logs. For each source user/IP address, it includes the correlation object that triggered the match and the match count, which is aggregated from the match evidence collated in the correlated events logs.For details see Use the Automated Correlation Engine.<br>Available on the PA-3000 Series, PA-5000 Series, PA-7000 Series, and Panorama.<br>Sort attributes: severity (by default) |
| **Hosts Visiting Malicious URLs** | Displays the frequency with which hosts (IP address/hostnames) on your network have accessed malicious URLs. These URLs are known to be malware based on categorization in PAN-DB.<br>Sort attributes: count<br>Charts available: line |

 © Palo Alto Networks, Inc.

| Widget | Description |
|---|---|
| **Hosts Resolving Malicious Domains** | Displays the top hosts matching DNS signatures; hosts on the network that are attempting to resolve the hostname or domain of a malicious URL. This information is gathered from an analysis of the DNS activity on your network. It utilizes passive DNS monitoring, DNS traffic generated on the network, activity seen in the sandbox if you have configured DNS sinkhole on the firewall, and DNS reports on malicious DNS sources that are available to Palo Alto Networks customers.<br>Sort attributes: count<br>Charts available: line |
| **Threat Activity** | Displays the threats seen on your network. This information is based on signature matches in Antivirus, Anti-Spyware, and Vulnerability Protection profiles and viruses reported by WildFire.<br>Sort attributes: threats<br>Charts available: bar, area, column |
| **WildFire Activity by Application** | Displays the applications that generated the most WildFire submissions. This widget uses the malicious and benign verdict from the WildFire Submissions log.<br>Sort attributes: malicious, benign<br>Charts available: bar, line |
| **WildFire Activity by File Type** | Displays the threat vector by file type. This widget displays the file types that generated the most WildFire submissions and uses the malicious and benign verdict from the WildFire Submissions log. If this data is unavailable, the widget is empty.<br>Sort attributes: malicious, benign<br>Charts available: bar, line |
| **Applications using Non Standard Ports** | Displays the applications that are entering your network on non-standard ports. If you have migrated your firewall rules from a port-based firewall, use this information to craft policy rules that allow traffic only on the default port for the application. Where needed, make an exception to allow traffic on a non-standard port or create a custom application.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: treemap, line |
| **Rules Allowing Applications On Non Standard Ports** | Displays the security policy rules that allow applications on non-default ports. The graph displays all the rules, while the table displays the top ten rules and aggregates the data from the remaining rules as other.<br>This information helps you identify gaps in network security by allowing you to assess whether an application is hopping ports or sneaking into your network. For example, you can validate whether you have a rule that allows traffic on any port except the default port for the application. Say for example, you have a rule that allow DNS traffic on its *application-default* port (port 53 is the standard port for DNS). This widget will display any rule that allows DNS traffic into your network on any port except port 53.<br>Sort attributes: bytes, sessions, threats, content, URLs<br>Charts available: treemap, line |

**Blocked Activity**—Focuses on traffic that was prevented from coming into the network

| Widget | Description |
|---|---|
| **Blocked Application Activity** | Displays the applications that were denied on your network, and allows you to view the threats, content, and URLs that you kept out of your network.<br>Sort attributes: threats, content, URLs<br>Charts available: treemap, area, column |
| **Blocked User Activity** | Displays user requests that were blocked by a match on an antivirus, anti-spyware, file blocking or url filtering profile attached to security policy.<br>Sort attributes: threats, content, URLs<br>Charts available: bar, area, column |
| **Blocked Threats** | Displays the threats that were successfully denied on your network. These threats were matched on antivirus signatures, vulnerability signatures, and DNS signatures available through the dynamic content updates on the firewall.<br>Sort attributes: threats<br>Charts available: bar, area, column |
| **Blocked Content** | Displays the files and data that was blocked from entering the network. The content was blocked because security policy denied access based on criteria defined in a File Blocking security profile or a Data Filtering security profile.<br>Sort attributes: files, data<br>Charts available: bar, area, column |
| **Security Policies Blocking Activity** | Displays the security policy rules that blocked or restricted traffic into your network. Because this widget displays the threats, content, and URLs that were denied access into your network, you can use it to assess the effectiveness of your policy rules. This widget does not display traffic that blocked because of deny rules that you have defined in policy.<br>Sort attributes: threats, content, URLs<br>Charts available: bar, area, column |

 © Palo Alto Networks, Inc.

## ACC Filters

The graphs and tables on the ACC widgets allow you to use filters to narrow the scope of data that is displayed, so that you can isolate specific attributes and analyze information you want to view in greater detail. The ACC supports the simultaneous use of widget and global filters.

- **Widget Filters**—Apply a widget filter, which is a filter that is *local* to a specific widget. A widget filter allows you to interact with the graph and customize the display so that you can drill down in to the details and access the information you want to monitor on a specific widget. To create a widget filter that is persistent across reboots, you must use the **Set Local Filter** option.

- **Global filters**—Apply global filters across all the tabs in the ACC. A global filter allows you to pivot the display around the details you care about right now and exclude the unrelated information from the current display. For example, to view all events relating to a specific user and application, you can apply the username and the application as a global filter and view only information pertaining to the user and the application through all the tabs and widgets on the ACC. Global filters are not persistent.

You can apply global filters in three ways:

- **Set a global filter from a table**—Select an attribute from a table in any widget and apply the attribute as a global filter.
- **Promote a widget filter to be a global filter**—Promote a value in a table or a graph to a global filter by using the dropdown menu next to the value. This option allows you to elevate a local filter used in a widget, and apply the attribute globally to update the display across all the tabs on the ACC.
- **Define a global filter**—Define a filter using the **Global Filters** pane on the ACC.

See Interact with the ACC for details on using these filters.

 © Palo Alto Networks, Inc.

## Interact with the ACC

To customize and refine the ACC display, you can add and delete tabs, add and delete widgets, set local and global filters, and interact with the widgets.

| Work with the Tabs and Widgets | |
|---|---|
| • Add a tab. | 1. Select the + icon along the list of tabs.<br>2. Add a **View Name**. This name will be used as the name for the tab. You can add up to five tabs. |
| • Edit a tab. | Select the tab, and click the pencil icon next to the tab name, to edit the tab. For example .<br>Editing a tab allow you to add or delete or reset the widgets that are displayed in the tab. You can also change the widget layout in the tab. |
| • See what widgets are included in a tab. | 1. Select the tab, and click on the pencil icon to edit it.<br>2. Select the **Add Widgets** drop-down and verify the widgets that have the check boxes selected. |
| • Add a widget or a widget group. | 1. Add a new tab or edit a predefined tab.<br>2. Select **Add Widget**, and then select the check box that corresponds to the widget you want to add. You can select up to a maximum of 12 widgets.<br>3. (Optional) To create a 2-column layout, select **Add Widget Group**. You can drag and drop widgets into the 2-column display. As you drag the widget into the layout, a placeholder will display for you to drop the widget.<br> You cannot name a widget group. |
| • Delete a tab or a widget group/ widget. | 1. To delete a custom tab, select the tab and click the X icon.<br> 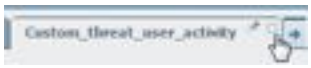 <br> You cannot delete a predefined tab.<br>2. To delete a widget group/widget, edit the tab and in the workspace section, click the [X] icon on the right. You cannot undo a deletion. |
| • Reset the default widgets in a tab. | On a predefined tab, such as the **Blocked Activity** tab, you can delete one or more widgets. If you want to reset the layout to include the default set of widgets for the tab, edit the tab and click **Reset View**. |
| • Zoom in on the details in an area, column, or line graph.<br>Watch how the zoom-in capability works. | Click and drag an area in the graph to zoom in. For example, when you zoom into a line graph, it triggers a re-query and the firewall fetches the data for the selected time period. It is not a mere magnification. |

| Work with the Tabs and Widgets (Continued) | |
|---|---|
| • Use the table drop-down to find more information on an attribute.<br> | 1. Hover over an attribute in a table to see the drop-down.<br>2. Click into the drop-down to view the available options.<br>• **Global Find**—Global Find allows you to find references to the attribute (username/IP address, object name, policy rule name, threat ID, or application name) anywhere in the candidate configuration on the firewall.<br>• **Promote Filter**—Sets the attribute as a global filter. This allows you to filter all the views in the ACC based on the filters you have applied.<br>• **Value**—Displays the details of the threat ID, or application name, or address object.<br>• **Who Is**—Performs a domain name (*WHOIS*) lookup for the IP address. The lookup queries databases that store the registered users or assignees of an Internet resource.<br>• **Search HIP Report**—Uses the username or IP address to find matches in a HIP Match report. |
| • Set a widget filter.<br>You can also click an attribute in the table (below the graph) to apply it as a widget filter. | 1. Select a widget and click the icon.<br>2. Click the icon to add the filters you want to apply.<br>3. Click **Apply**. These filters are persistent across reboots.<br>The active widget filters are indicated next to the widget name. |
| • Negate a widget filter | 1. Click the icon to display the Setup Local Filters dialog.<br>2. Add a filter, and then click the negate icon. |
| • Set a global filter from a table. | 1. Hover over an attribute in the table below the chart, and click the drop-down.<br>2. Click **Promote Filter** to add the attribute as a global filter.<br> 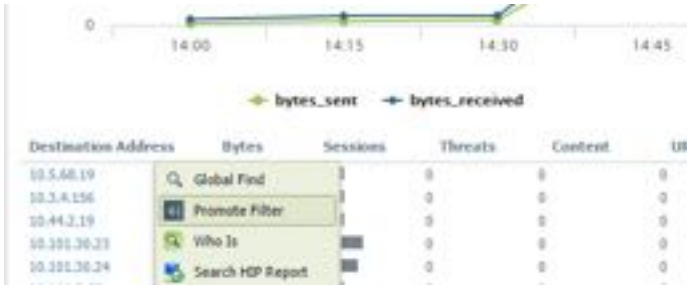 |
| • Set a global filter using the Global Filters pane.<br>Watch global filters in action. | 1. Locate the **Global Filters** pane on the left side of the ACC.<br> 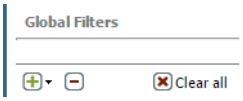 <br>2. Click the icon to view the list of filters you can apply. |

 © Palo Alto Networks, Inc.

| Work with the Tabs and Widgets (Continued) | |
|---|---|
| • Promote a widget filter to a global filter. | 1. On any table in a widget, click the link for an attribute. This sets the attribute as a widget filter.<br>2. To promote the filter to be a global filter, select the arrow to the right of the filter.<br> 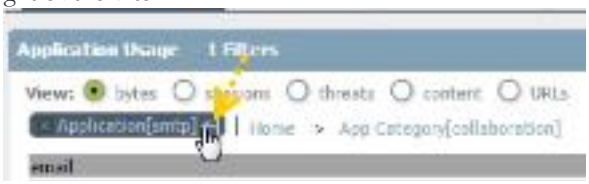 |
| • Remove a filter. | Click the icon to remove a filter.<br>• For global filters: It is located in the Global Filters pane.<br>• For widget filters: Click the icon to display the Setup Local Filters dialog, then select the filter, and click the icon.<br> |
| • Clear all filters. | • For global filters: Click the **Clear All** button under Global Filters.<br>• For widget filters: Select a widget and click the icon. Then click the **Clear All** button in the Setup Local Filters dialog. |
| • See what filters are in use. | • For global filters: The number of global filters applied are displayed on the left pane under Global Filters.<br>• For widget filters: The number of widget filters applied on a widget are displayed next to the widget name. To view the filters, click the icon. |
| • Reset the display on a widget. | • If you set a widget filter or drill into a graph, click the **Home** link to reset the display in the widget.<br> |

## Use Case: ACC—Path of Information Discovery

The ACC has a wealth of information that you can use as a starting point for analyzing network traffic. Let's look at an example on using the ACC to uncover events of interest. This example illustrates how you can use the ACC to ensure that legitimate users can be held accountable for their actions, detect and track unauthorized activity, and detect and diagnose compromised hosts and vulnerable systems on your network.

The widgets and filters in the ACC give you the capability to analyze the data and filter the views based on events of interest or concern. You can trace events that pique your interest, directly export a PDF of a tab, access the raw logs, and save a personalized view of the activity that you want to track. These capabilities make it possible for you to monitor activity and develop policies and countermeasures for fortifying your network against malicious activity. In this section, you will Interact with the ACC widgets across different tabs, drill down using widget filters, and pivot the ACC views using global filters, and export a PDF for sharing with incidence response or IT teams.

At first glance, you see the Application Usage and User Activity widgets in the **ACC > Network Activity** tab. The User Activity widget shows that user Marsha Wirth has transferred 718 Megabytes of data during the last hour. This volume is nearly six times more than any other user on the network. To see the trend over the past few hours, expand the **Time** period to the **Last 6 Hrs**, and now Marsha's activity has been 6.5 Gigabytes over 891 sessions and has triggered 38 threats signatures.

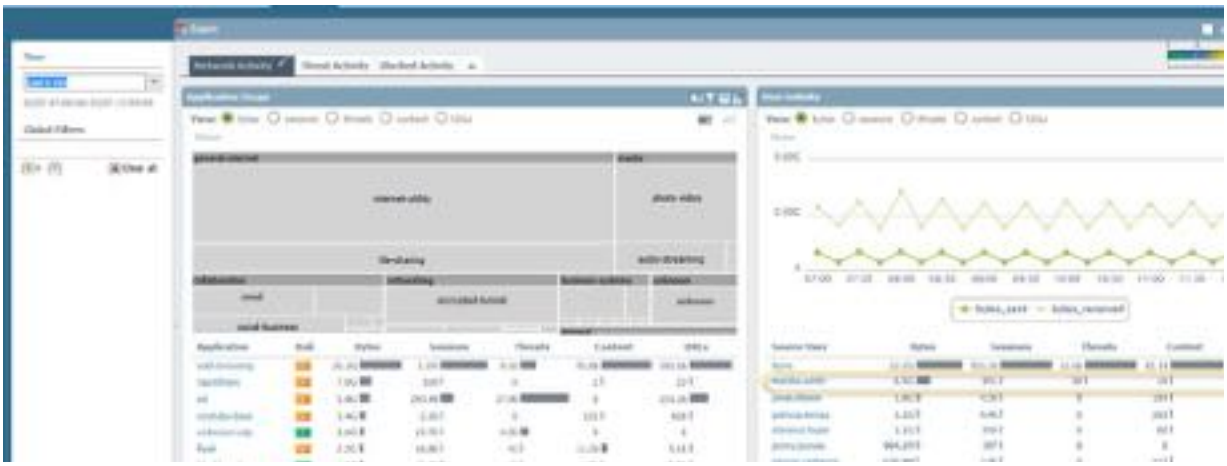

Because Marsha has transferred a large volume of data, apply her username as a global filter (ACC Filters) and pivot all the views in the ACC to Marsha's traffic activity.

 © Palo Alto Networks, Inc.

The Application Usage tab now shows that the top application that Martha used was rapidshare, a Swiss-owned file-hosting site that belongs to the file-sharing URL category. For further investigation, add rapidshare as a global filter, and view Marsha's activity in the context of rapidshare.

Consider whether you want to sanction rapidshare for company use. Should you allow uploads to this site and do you need a QoS policy to limit bandwidth?

To view which IP addresses Marsha has communicated with, check the **Destination IP Activity** widget, and view the data by bytes and by URLs.

To know which countries Marsha communicated with, sort on **sessions** in the **Destination Regions** widget.

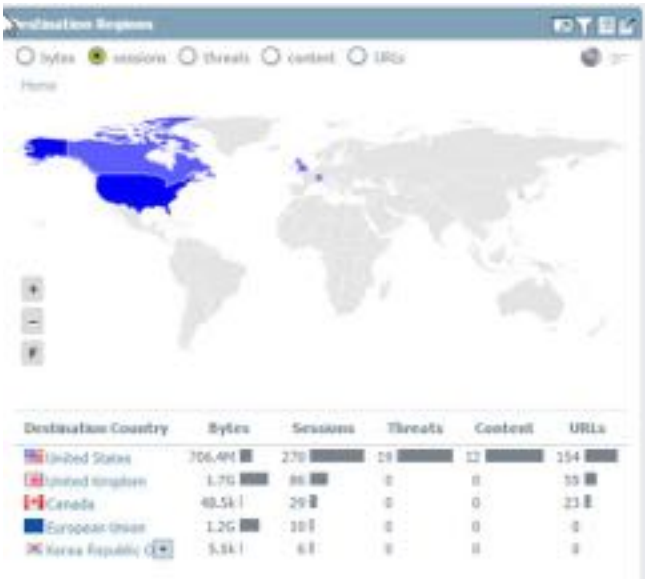


From this data, you can confirm that Marsha, a user on your network, has established sessions in Korea and the European Union, and she logged 19 threats in her sessions within the United States.

To look at Marsha's activity from a threat perspective, remove the global filter for rapidshare. In the **Threat Activity** widget on the **Threat Activity** tab, view the threats. The widget displays that her activity had triggered a match for 26 vulnerabilities in the overflow, DoS and code-execution threat category. Several of these vulnerabilities are of critical severity.

To further drill-down into each vulnerability, click into the graph and narrow the scope of your investigation. Each click automatically applies a local filter on the widget.

 © Palo Alto Networks, Inc.

To investigate each threat by name, you can create a global filter for say, **Microsoft Works File Converter Field Length Remote Code Execution Vulnerability**. Then, view the **User Activity widget** in the **Network Activity** tab. The tab is automatically filtered to display threat activity for Marsha (notice the global filters in the screenshot).

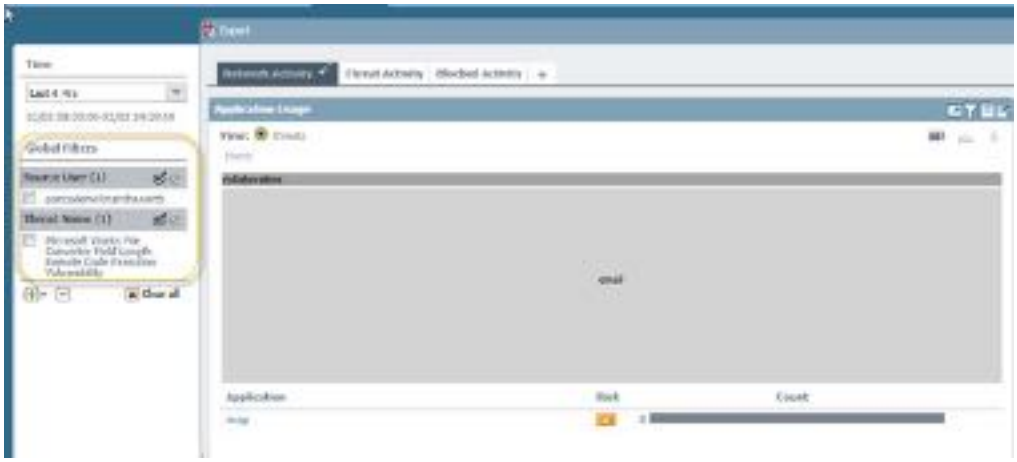

Notice that this Microsoft code-execution vulnerability was triggered over email, by the imap application. You can now establish that Martha has IE vulnerabilities and email attachment vulnerabilities, and perhaps her computer needs to be patched. You can now either navigate to the **Blocked Threats** widget in the **Blocked Activity** tab to check how many of these vulnerabilities were blocked.

Or, you can check the **Rule Usage** widget on the **Network Activity** tab to discover how many vulnerabilities made it into your network and which security rule allowed this traffic, and navigate directly to the security rule using the **Global Find** capability.

Then, drill into why imap used a non-standard port 43206 instead of port 143, which is the default port for the application. Consider modifying the security policy rule to allow applications to only use the default port for the application, or assess whether this port should be an exception on your network.

To review if any threats were logged over imap, check Marsha's activity in the **WildFire Activity by Application** widget in the **Threat Activity** tab. You can confirm that Marsha had no malicious activity, but to verify that other no other user was compromised by the imap application, negate Marsha as a global filter and look for other users who triggered threats over imap.

Click into the bar for imap in the graph and drill into the inbound threats associated with the application. To find out who an IP address is registered to, hover over the attacker IP address and select the **Who Is** link in the drop-down.

Because the session count from this IP address is high, check the **Blocked Content** and **Blocked Threats** widgets in the **Blocked Activity** tab for events related to this IP address. The **Blocked Activity** tab allows you to validate whether or not your policy rules are effective in blocking content or threats when a host on your network is compromised.

 © Palo Alto Networks, Inc.

Use the **Export PDF** capability on the ACC to export the current view (create a snapshot of the data) and send it to an incidence response team. To view the threat logs directly from the widget, you can also click the icon to jump to the logs; the query is generated automatically and only the relevant logs are displayed onscreen (for example in **Monitor > Logs > Threat Logs**).

You have now used the ACC to review network data/trends to find which applications or users are generating the most traffic, and how many application are responsible for the threats seen on the network. You were able to identify which application(s), user(s) generated the traffic, determine whether the application was on the default port, and which policy rule(s) allowed the traffic into the network, and determine whether the threat is spreading laterally on the network. You also identified the destination IP addresses, geo-locations with which hosts on the network are communicating with. Use the conclusions from your investigation to craft goal-oriented policies that can secure users and your network.

# App Scope

The App Scope reports provide visibility and analysis tools to help pinpoint problematic behavior, helping you understand changes in application usage and user activity, users and applications that take up most of the network bandwidth, and identify network threats.

With the App Scope reports, you can quickly see if any behavior is unusual or unexpected. Each report provides a dynamic, user-customizable window into the network; hovering the mouse over and clicking either the lines or bars on the charts opens detailed information about the specific application, application category, user, or source on the ACC. The App Scope charts on **Monitor > App Scope** give you the ability to:

- Toggle the attributes in the legend to only view chart details that you want to review. The ability to include or exclude data from the chart allows you to change the scale and review details more closely.
- Click into an attribute in a bar chart and drill down to the related sessions in the ACC. Click into an Application name, Application Category, Threat Name, Threat Category, Source IP address or Destination IP address on any bar chart to filter on the attribute and view the related sessions in the ACC.
- Export a chart or map to PDF or as an image. For portability and offline viewing, you can Export charts and maps as PDFs or PNG images.

The following App Scope reports are available:

- Summary Report
- Change Monitor Report
- Threat Monitor Report
- Threat Map Report
- Network Monitor Report
- Traffic Map Report

 © Palo Alto Networks, Inc.

## Summary Report

The App Scope Summary report (**Monitor > App Scope > Summary**) displays charts for the top five gainers, losers, and bandwidth consuming applications, application categories, users, and sources.

## Change Monitor Report

The App Scope Change Monitor report (**Monitor > App Scope > Change Monitor**) displays changes over a specified time period. For example, the following chart displays the top applications that gained in use over the last hour as compared with the last 24-hour period. The top applications are determined by session count and sorted by percent.

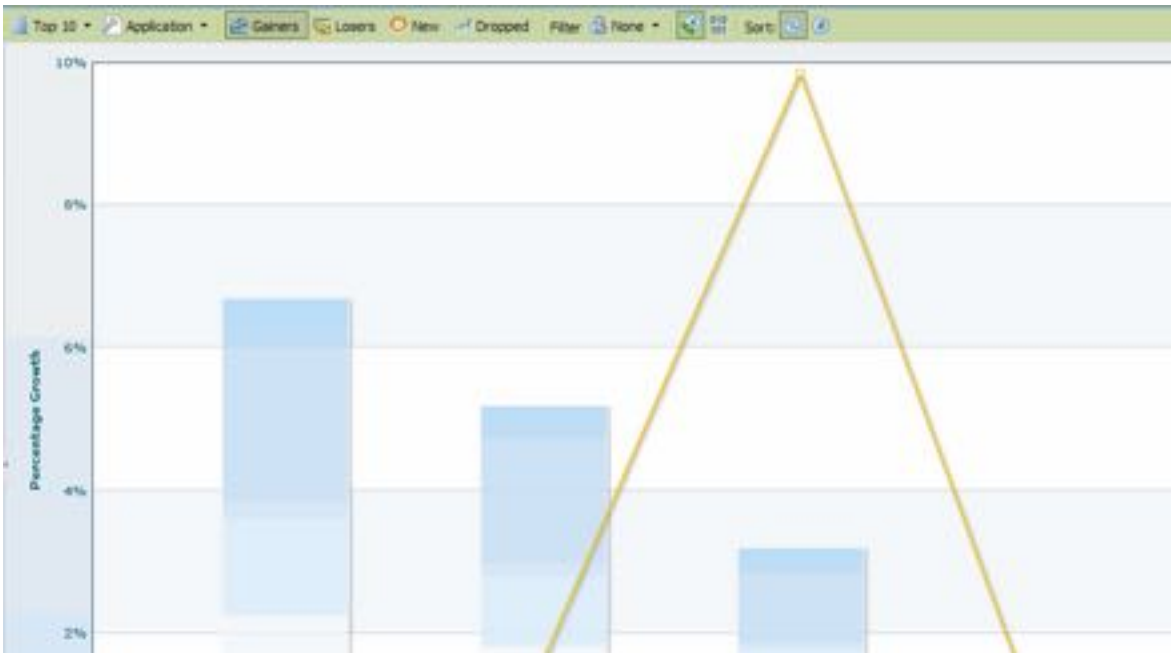


The Change Monitor Report contains the following buttons and options.

| Button | Description |
|---|---|
| Top 10 | Determines the number of records with the highest measurement included in the chart. |
| Application | Determines the type of item reported: Application, Application Category, Source, or Destination. |
| Gainers | Displays measurements of items that have increased over the measured period. |
| Losers | Displays measurements of items that have decreased over the measured period. |
| New | Displays measurements of items that were added over the measured period. |
| Dropped | Displays measurements of items that were discontinued over the measured period. |

 © Palo Alto Networks, Inc.

| Button | Description |
|---|---|
| 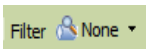 | Applies a filter to display only the selected item. **None** displays all entries. |
| | Determines whether to display session or byte information. |
| 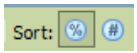 | Determines whether to sort entries by percentage or raw growth. |
| 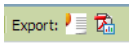 | Exports the graph as a .png image or as a PDF. |
| 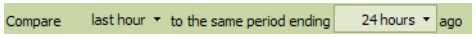 | Specifies the period over which the change measurements are taken. |

## Threat Monitor Report

The App Scope Threat Monitor report (**Monitor > App Scope > Threat Monitor**) displays a count of the top threats over the selected time period. For example, the following figure shows the top 10 threat types over the last 6 hours.

Each threat type is color-coded as indicated in the legend below the chart. The Threat Monitor report contains the following buttons and options.

| Button | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| 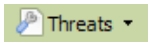 | Determines the type of item measured: Threat, Threat Category, Source, or Destination. |
| | Applies a filter to display only the selected type of items. |
| | Determines whether the information is presented in a stacked column chart or a stacked area chart. |
| 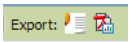 | Exports the graph as a .png image or as a PDF. |
| | Specifies the period over which the measurements are taken. |

 © Palo Alto Networks, Inc.

## Threat Map Report

The App Scope Threat Map report (**Monitor > App Scope > Threat Map**) shows a geographical view of threats, including severity. Each threat type is color-coded as indicated in the legend below the chart.

The firewall uses geolocation for creating threat maps. The firewall is placed at the bottom of the threat map screen, if you have not specified the geolocation coordinates (**Device > Setup > Management,** General Settings section) on the firewall.

The Threat Map report contains the following buttons and options.

| Button | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| 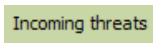 | Displays incoming threats. |
| 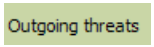 | Displays outgoing threats. |
| | Applies a filter to display only the selected type of items. |
|  | Zoom in and zoom out of the map. |
| 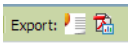 | Exports the graph as a .png image or as a PDF. |
| 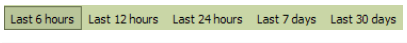 | Indicates the period over which the measurements are taken. |

## Network Monitor Report

The App Scope Network Monitor report (**Monitor > App Scope > Network Monitor**) displays the bandwidth dedicated to different network functions over the specified period of time. Each network function is color-coded as indicated in the legend below the chart. For example, the image below shows application bandwidth for the past 7 days based on session information.

The Network Monitor report contains the following buttons and options.

| Button | Description |
|---|---|
| | Determines the number of records with the highest measurement included in the chart. |
| 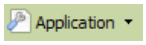 | Determines the type of item reported: Application, Application Category, Source, or Destination. |
| 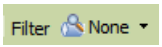 | Applies a filter to display only the selected item. **None** displays all entries. |
| | Determines whether to display session or byte information. |
| 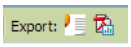 | Exports the graph as a .png image or as a PDF. |
| | Determines whether the information is presented in a stacked column chart or a stacked area chart. |
| | Indicates the period over which the change measurements are taken. |

 © Palo Alto Networks, Inc.

## Traffic Map Report

The App Scope Traffic Map (**Monitor > App Scope > Traffic Map**) report shows a geographical view of traffic flows according to sessions or flows.

The firewall uses geolocation for creating traffic maps. The firewall is placed at the bottom of the traffic map screen, if you have not specified the geolocation coordinates (**Device > Setup > Management,** General Settings section) on the firewall.

Each traffic type is color-coded as indicated in the legend below the chart. The Traffic Map report contains the following buttons and options.

| Buttons | Description |
|---|---|
| Top 10 | Determines the number of records with the highest measurement included in the chart. |
| Incoming threats | Displays incoming threats. |
| Outgoing threats | Displays outgoing threats. |
| 010 101 | Determines whether to display session or byte information. |
| Zoom In Zoom Out | Zoom in and zoom out of the map. |
| Export: | Exports the graph as a .png image or as a PDF. |
| Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days | Indicates the period over which the change measurements are taken. |

# Use the Automated Correlation Engine

The automated correlation engine is an analytics tool that uses the logs on the firewall to detect actionable events on your network. The engine correlates a series of related threat events that, when combined, indicate a likely compromised host on your network or some other higher level conclusion. It pinpoints areas of risk, such as compromised hosts on the network, allows you to assess the risk and take action to prevent exploitation of network resources. The automated correlation engine uses *correlation objects* to analyze the logs for patterns and when a match occurs, it generates a *correlated event*.

The automated correlation engine is supported on the following platforms:

- Panorama—M-Series appliance and the virtual appliance
- PA-7000 Series firewall
- PA-5000 Series firewall
- PA-3000 Series firewall


- Automated Correlation Engine Concepts
- View the Correlated Objects
- Interpret Correlated Events
- Use the Compromised Hosts Widget in the ACC


 © Palo Alto Networks, Inc.

# Automated Correlation Engine Concepts

The automated correlation engine uses *correlation objects* to analyze the logs for patterns and when a match occurs, it generates a *correlated event*.

- Correlation Object
- Correlated Events

## Correlation Object

A correlation object is a definition file that specifies patterns to match against, the data sources to use for the lookups, and time period within which to look for these patterns. A pattern is a boolean structure of conditions that queries the following data sources (or logs) on the firewall: application statistics, traffic, traffic summary, threat summary, threat, data filtering, and URL filtering. Each pattern has a severity rating, and a threshold for the number of times the pattern match must occur within a defined time limit to indicate malicious activity. When the match conditions are met, a correlated event is logged.

A correlation object can connect isolated network events and look for patterns that indicate a more significant event. These objects identify suspicious traffic patterns and network anomalies, including suspicious IP activity, known command-and-control activity, known vulnerability exploits, or botnet activity that, when correlated, indicate with a high probability that a host on the network has been compromised. Correlation objects are defined and developed by the Palo Alto Networks Threat Research team, and are delivered with the weekly dynamic updates to the firewall and Panorama. To obtain new correlation objects, the firewall must have a Threat Prevention license. Panorama requires a support license to get the updates.

The patterns defined in a correlation object can be static or dynamic. Correlated objects that include patterns observed in WildFire are dynamic, and can correlate malware patterns detected by WildFire with command-and-control activity initiated by a host that was targeted with the malware on your network. For example, when a host submits a file to the WildFire cloud and the verdict is malicious, the correlation object looks for other hosts or clients on the network that exhibit the same behavior seen in the cloud. If the malware sample had performed a DNS query and browsed to a malware domain, the correlation object will parse the logs for a similar event. When the activity on a host matches the analysis in the cloud, a high severity correlated event is logged.

## Correlated Events

A correlated event is logged when the patterns and thresholds defined in a correlation object match the traffic patterns on your network. To Interpret Correlated Events and to view a graphical display of the events, see Use the Compromised Hosts Widget in the ACC.

# View the Correlated Objects

**View the Correlation Objects Available on the Firewall**

Step 1 To view the correlation objects that are currently available, select **Monitor > Automated Correlation Engine > Correlation Objects**. All the objects in the list are enabled by default.

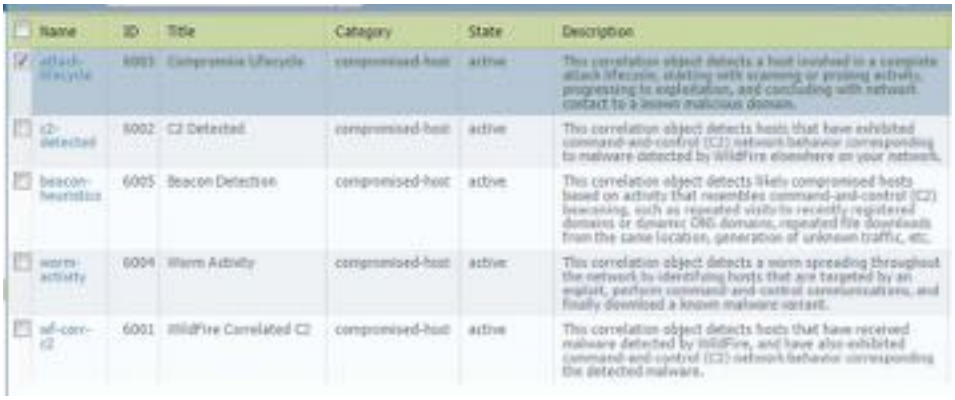


Step 2 View the details on each correlation object. Each object provides the following information:

- **Name** and **Title**—The name and title indicate the type of activity that the correlation object detects. The name column is hidden from view, by default. To view the definition of the object, unhide the column and click the name link.
- **ID**—A unique number that identifies the correlation object; this column is also hidden by default. The IDs are in the 6000 series.
- **Category**—A classification of the kind of threat or harm posed to the network, user, or host. For now, all the objects identify compromised hosts on the network.
- **State**—Indicates whether the correlation object is enabled (active) or disabled (inactive). All the objects in the list are enabled by default, and are hence active. Because these objects are based on threat intelligence data and are defined by the Palo Alto Networks Threat Research team, keep the objects active in order to track and detect malicious activity on your network.
- **Description**—Specifies the match conditions for which the firewall or Panorama will analyze logs. It describes the sequence of conditions that are matched on to identify acceleration or escalation of malicious activity or suspicious host behavior. For example, the **Compromise Lifecycle** object detects a host involved in a complete attack lifecycle in a three-step escalation that starts with scanning or probing activity, progressing to exploitation, and concluding with network contact to a known malicious domain.

For more information, see Automated Correlation Engine Concepts and Use the Automated Correlation Engine.

 © Palo Alto Networks, Inc.

## Interpret Correlated Events

You can view and analyze the logs generated for each correlated event in the **Monitor > Automated Correlation Engine > Correlated Events** tab.

Correlated Events includes the following details:

| Field | Description |
|---|---|
| **Match Time** | The time the correlation object triggered a match. |
| **Update Time** | The time when the event was last updated with evidence on the match. As the firewall collects evidence on pattern or sequence of events defined in a correlation object, the time stamp on the correlated event log is updated. |
| **Object Name** | The name of the correlation object that triggered the match. |
| **Source Address** | The IP address of the user/device on your network from whom.which the traffic originated. |
| **Source User** | The user and user group information from the directory server, if User-ID is enabled. |

| Field | Description |
|---|---|
| **Severity**<br> To configure the firewall or Panorama to send alerts using email, SNMP or syslog messages for a desired severity level, see Use External Services for Monitoring. | A rating that indicates the urgency and impact of the match. The severity level indicates the extent of damage or escalation pattern, and the frequency of occurrence. Because correlation objects are primarily for detecting threats, the correlated events typically relate to identifying compromised hosts on the network and the severity implies the following:<br>• **Critical**—Confirms that a host has been compromised based on correlated events that indicate an escalation pattern. For example, a critical event is logged when a host that received a file with a malicious verdict by WildFire exhibits the same command-and-control activity that was observed in the WildFire sandbox for that malicious file.<br>• **High**—Indicates that a host is very likely compromised based on a correlation between multiple threat events, such as malware detected anywhere on the network that matches the command-and-control activity generated by a particular host.<br>• **Medium**—Indicates that a host is likely compromised based on the detection of one or multiple suspicious events, such as repeated visits to known malicious URLs, which suggests a scripted command-and-control activity.<br>• **Low**—Indicates that a host is possibly compromised based on the detection of one or multiple suspicious events, such as a visit to a malicious URL or a dynamic DNS domain.<br>• **Informational**—Detects an event that may be useful in aggregate for identifying suspicious activity, but the event is not necessarily significant on its own. |
| **Summary** | A description that summarizes the evidence gathered on the correlated event. |

 © Palo Alto Networks, Inc.

Click the icon to see the detailed log view, which includes all the evidence on a match:

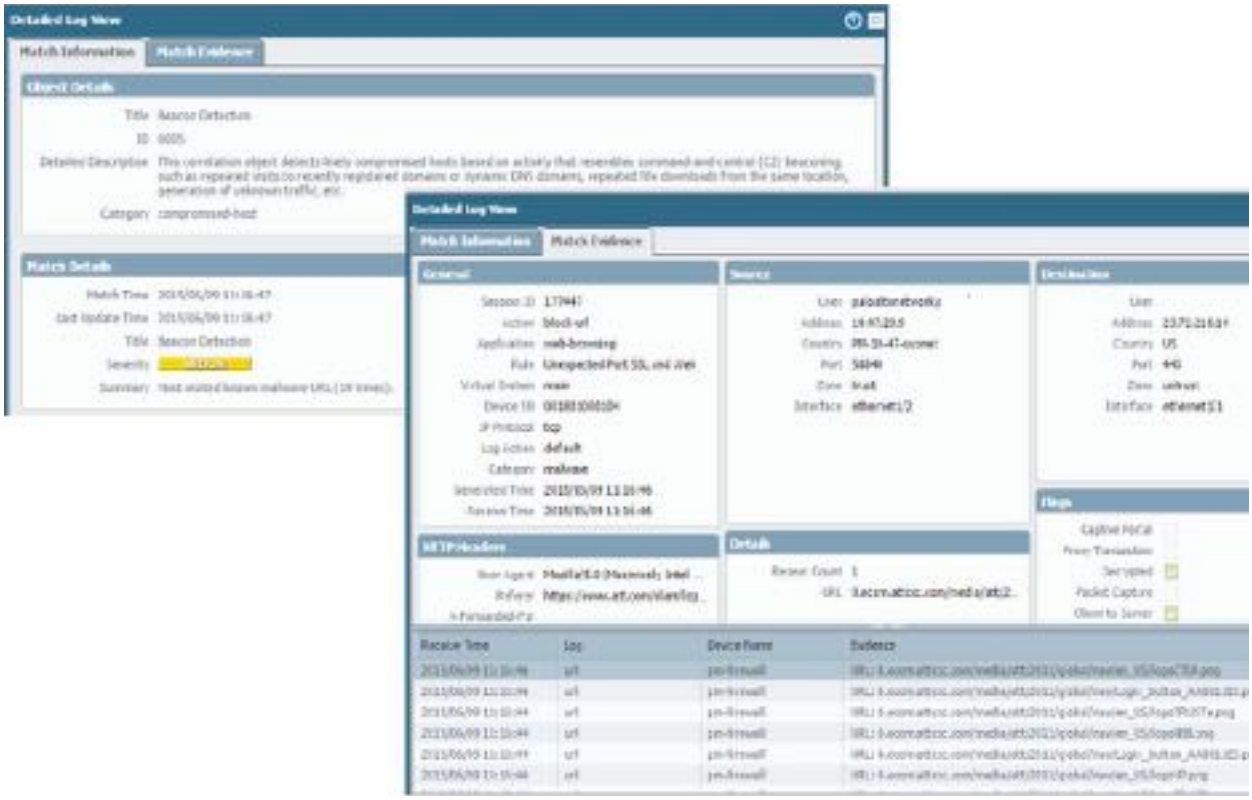

| Tab | Description |
|---|---|
| **Match Information** | Object Details: Presents information on the Correlation Object that triggered the match. |
| | Match Details: A summary of the match details that includes the match time, last update time on the match evidence, severity of the event, and an event summary. |
| **Match Evidence** | Presents all the evidence that corroborates the correlated event. It lists detailed information on the evidence collected for each session. |

## Use the Compromised Hosts Widget in the ACC

The compromised hosts widget on **ACC >Threat Activity**, aggregates the Correlated Events and sorts them by severity. It displays the source IP address/user who triggered the event, the correlation object that was matched and the number of times the object was matched. Use the match count link to jump to the match evidence details.

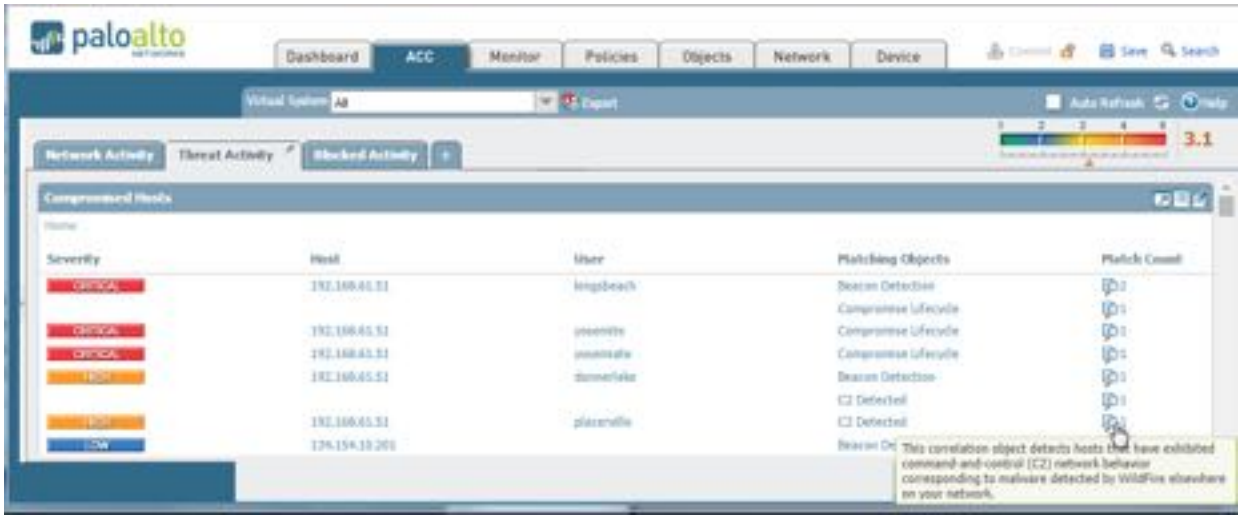


For more details, see Use the Automated Correlation Engine and Use the Application Command Center.

 © Palo Alto Networks, Inc.

# Take Packet Captures

All Palo Alto Networks firewalls allow you to take packet captures (pcaps) of traffic that traverses the management interface and network interfaces on the firewall. When taking packet captures on the dataplane, you may need to Disable Hardware Offload to ensure that the firewall captures all traffic.

Packet capture can be very CPU intensive and can degrade firewall performance. Only use this feature when necessary and make sure you turn it off after you have collected the required packets.

There are four different types of packet captures you can enable, depending on what you need to do:

- **Custom Packet Capture**—The firewall captures packets for all traffic or for specific traffic based on filters that you define. For example, you can configure the firewall to only capture packets to and from a specific source and destination IP address or port. You then use the packet captures for troubleshooting network-related issues or for gathering application attributes to enable you to write custom application signatures or to request an application signature from Palo Alto Networks. See Take a Custom Packet Capture.
- **Threat Packet Capture**—The firewall captures packets when it detects a virus, spyware, or vulnerability. You enable this feature in Antivirus, Anti-Spyware, and Vulnerability Protection security profiles. A link to view or export the packet captures will appear in the second column of the Threat log. These packet captures provide context around a threat to help you determine if an attack is successful or to learn more about the methods used by an attacker. You can also submit this type of pcap to Palo Alto Networks to have a threat re-analyzed if you feel its a false-positive or false-negative. See Take a Threat Packet Capture.
- **Application Packet Capture**—The firewall captures packets based on a specific application and filters that you define. A link to view or export the packet captures will appear in the second column of the Traffic logs for traffic that matches the packet capture rule. See Take an Application Packet Capture.
- **Management Interface Packet Capture**—The firewall captures packets on the management interface (MGT) The packet captures are useful when troubleshooting services that traverse the interface, such as device management authentication to external servers (LDAP and RADIUS for example), software and content updates, log forwarding, communication with SNMP servers, and authentication requests for GlobalProtect and Captive Portal. See Take a Packet Capture on the Management Interface.

© Palo Alto Networks, Inc.

## Disable Hardware Offload

Packet captures on a Palo Alto Networks firewall are performed in the dataplane CPU, unless you configure the firewall to Take a Packet Capture on the Management Interface, in which case the packet capture is performed on the management plane. When a packet capture is performed on the dataplane, during the ingress stage, the firewall performs packet parsing checks and discards any packets that do not match the packet capture filter. Any traffic that is offloaded to the field-programmable gate array (FPGA) offload processor is also excluded, unless you turn off hardware offload. For example, encrypted traffic (SSL/SSH), network protocols (OSPF, BGP, RIP), application overrides, and terminating applications can be offloaded to the FPGA and therefore are excluded from packet captures by default. Some types of sessions will never be offloaded, such as ARP, all non-IP traffic, IPSec, VPN sessions, SYN, FIN, and RST packets.

Hardware offload is supported on the following firewalls: PA-2000 Series, PA-3050, PA-4000 Series, PA-5000 Series, and PA-7000 Series firewall.

Disabling hardware offload increases the dataplane CPU usage. If dataplane CPU usage is already high, you may want to schedule a maintenance window before disabling hardware offload.

| **Enable/Disable Hardware Offload** | |
|---|---|
| Step 1 | Disable hardware offload by running the following CLI command:<br>`admin@PA-7050> set session offload no` |
| Step 2 | After the firewall captures the required traffic, enable hardware offload by running the following CLI command:<br>`admin@PA-7050> set session offload yes` |

## Take a Custom Packet Capture

Custom packet captures allow you to define the traffic that the firewall will capture. To ensure that you capture all traffic, you may need to Disable Hardware Offload.

 © Palo Alto Networks, Inc.

### Take a Custom Packet Capture

Step 1 Before you start a packet capture, identify the attributes of the traffic that you want to capture.

For example, to determine the source IP address, source NAT IP address, and the destination IP address for traffic between two systems, perform a ping from the source system to the to the destination system. After the ping is complete, go to **Monitor > Traffic** and locate the traffic log for the two systems. Click the **Detailed Log View** icon located in the first column of the log and note the source address, source NAT IP, and the destination address.

In the example that follows, we will use a packet capture to troubleshoot a Telnet connectivity issue from a user in the Trust zone to a server in the DMZ zone.

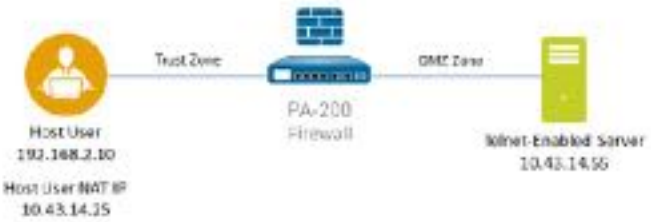

**Take a Custom Packet Capture (Continued)**

Step 2 Set packet capture filters, so the firewall only captures traffic you are interested in.

Filters will make it easier for you to locate the information you need in the packet capture and will reduce the processing power required by the firewall to take the packet capture. To capture all traffic, do not define filters and leave the filter option off.

For example, if you configured NAT on the firewall, you will need to apply two filters. The first one filters on the pre-NAT source IP address to the destination IP address and the second one filters traffic from the destination server to the source NAT IP address.

1. Select **Monitor > Packet Capture**.
2. Click **Clear All Settings** at the bottom of the window to clear any existing capture settings.
3. Click **Manage Filters** and click **Add**.
4. Select **Id** 1 and in the **Source** field enter the source IP address you are interested in and in the **Destination** field enter a destination IP address.

   For example, enter the source IP address `192.168.2.10` and the destination IP address `10.43.14.55`. To further filter the capture, set **Non-IP** to **exclude** non-IP traffic, such as broadcast traffic.
5. **Add** the second filter and select **Id** 2.

   For example, in the **Source** field enter `10.43.14.55` and in the **Destination** field enter `10.43.14.25`. In the **Non-IP** drop-down menu select **exclude**.

6. Click **OK**.

Step 3 Set **Filtering** to **On**.

 © Palo Alto Networks, Inc.

**Take a Custom Packet Capture (Continued)**

Step 4 Specify the traffic stage(s) that trigger the packet capture and the filename(s) to use to store the captured content. For a definition of each stage, click the **Help** icon on the packet capture page.

For example, to configure all packet capture stages and define a filename for each stage, perform the following procedure:

1. **Add** a **Stage** to the packet capture configuration and define a **File** name for the resulting packet capture.

   For example, select **receive** as the **Stage** and set the **File** name to `telnet-test-received`.

2. Continue to **Add** each **Stage** you want to capture (**firewall**, **transmit**, and **drop**) and set a unique **File** name for each stage.

Step 5 Set **Packet Capture** to **ON**.

Note the warning that system performance can be degraded and then click **OK**. If you define filters, the packet capture should have little impact on performance, but you should always turn **Off** packet capture after the firewall captures the data that you want to analyze.

**Take a Custom Packet Capture (Continued)**

Step 6 Generate traffic that matches the filters that you defined.

For this example, generate traffic from the source system to the Telnet-enabled server by running the following command from the source system (192.168.2.10):

```
telnet 10.43.14.55
```

Step 7 Turn packet capture **OFF** and then click the refresh icon to see the packet capture files.

Notice that in this case, there were no dropped packets, so the firewall did not create a file for the drop stage.

Step 8 Download the packet captures by clicking the filename in the File Name column.

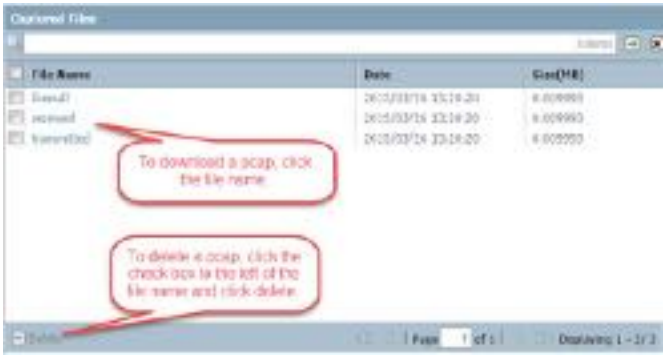


Step 9 View the packet capture files using a network packet analyzer, such as Wireshark.

In this example, the received.pcap packet capture shows a failed Telnet session from the source system at 192.168.2.10 to the Telnet-enabled server at 10.43.14.55. The source system sent the Telnet request to the server, but the server did not respond. In this example, the server may not have Telnet enabled, so check the server.

Step 10 Enable the Telnet service on the destination server (10.43.14.55) and turn on packet capture to take a new packet capture.

Step 11 Generate traffic that will trigger the packet capture.

Run the Telnet session again from the source system to the Telnet-enabled server

```
telnet 10.43.14.55
```

 © Palo Alto Networks, Inc.

**Take a Custom Packet Capture (Continued)**

Step 12 Download and open the received.pcap file and view it using a network packet analyzer.

The following packet capture now shows a successful Telnet session from the host user at 192.168.2.10 to the Telnet-enabled server at 10.43.14.55. Note that you also see the NAT address 10.43.14.25. When the server responds, it does so to the NAT address. You can see the session is successful as indicated by the three-way handshake between the host and the server and then you see Telnet data.

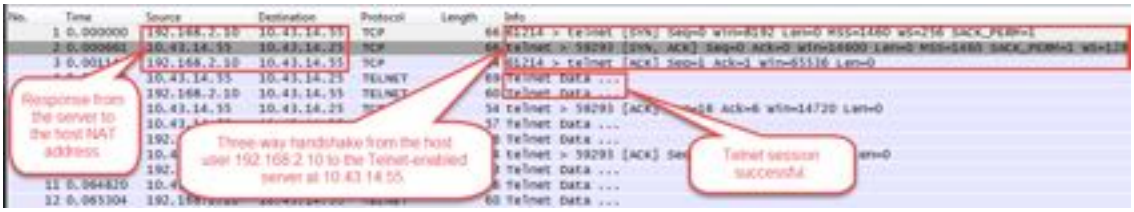


## Take a Threat Packet Capture

To configure the firewall to take a packet capture (pcap) when it detects a threat, enable packet capture on Antivirus, Anti-Spyware, and Vulnerability Protection security profiles.

**Take a Threat Packet Capture**

Step 1 Enable the packet capture option in the security profile.

Some security profiles allow you to define a single-packet capture, or extended-capture. If you choose extended-capture, define the capture length. This will allow the firewall to capture more packets to provide additional context related to the threat.

The firewall can only capture packets if the action for a given threat is set to allow or alert.

1. Select **Objects > Security Profiles** and enable the packet capture option for the supported profiles as follows:
   - **Antivirus**—Select a custom antivirus profile and in the **Antivirus** tab select the **Packet Capture** check box.
   - **Anti-Spyware**—Select a custom Anti-Spyware profile, click the **DNS Signatures** tab and in the **Packet Capture** drop-down, select **single-packet** or **extended-capture**.
   - **Vulnerability Protection**—Select a custom Vulnerability Protection profile and in the **Rules** tab, click **Add** to add a new rule, or select an existing rule. Set **Packet Capture** to **single-packet** or **extended-capture**. Note that if the profile has signature exceptions defined, click the **Exceptions** tab and in the **Packet Capture** column for a signature, set **single-packet** or **extended-capture**.
2. (Optional) If you selected **extended-capture** for any of the profiles, define the extended packet capture length.
   a. Select **Device > Setup > Content-ID** and edit the Content-ID Settings.
   b. In the **Extended Packet Capture Length (packets)** section, specify the number of packets that the firewall will capture (range is 1-50; default is 5).
   c. Click **OK**.

| Take a Threat Packet Capture (Continued) | | |
|---|---|---|
| Step 2 | Add the security profile (with packet capture enabled) to a Security Policy rule. | 1. Select **Policies > Security** and select a rule.<br>2. Select the **Actions** tab.<br>3. In the Profile Settings section, select a profile that has packet capture enabled.<br>For example, click the **Antivirus** drop-down and select a profile that has packet capture enabled. |

Step 3 View/export the packet capture from the Threat logs.

1. Select **Monitor > Logs > Threat**.
2. In the log entry that you are interested in, click the green packet capture icon in the second column. View the packet capture directly or **Export** it to your system.

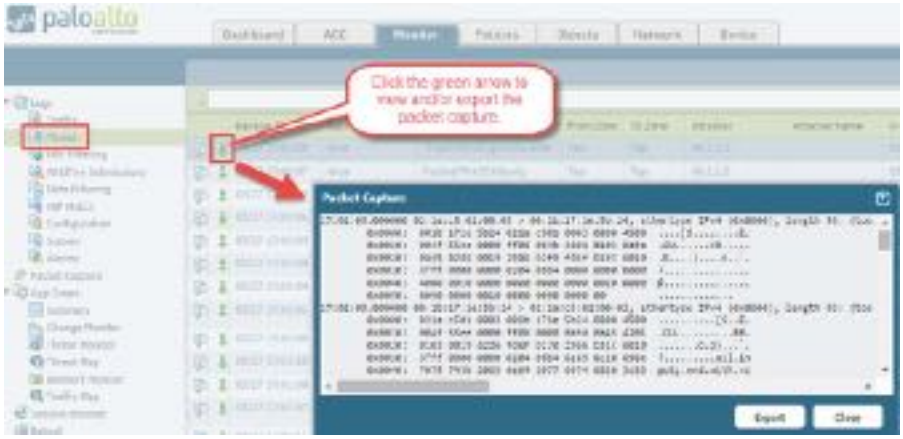


## Take an Application Packet Capture

The following topics describe two ways that you can configure the firewall to take application packet captures:

- Take a Packet Capture for Unknown Applications
- Take a Custom Application Packet Capture

### Take a Packet Capture for Unknown Applications

Palo Alto Networks firewalls automatically generate a packet capture for sessions that contain an application that it cannot identify. Typically, the only applications that are classified as unknown traffic—tcp, udp or non-syn-tcp—are commercially available applications that do not yet have App-ID signatures, are internal or custom applications on your network, or potential threats. You can use these packet captures to gather more context related to the unknown application or use the information to analyze the traffic for potential threats. You can also Manage Custom or Unknown Applications by controlling them through security policy or by

 © Palo Alto Networks, Inc.

writing a custom application signature and creating a security rule based on the custom signature. If the application is a commercial application, you can submit the packet capture to Palo Alto Networks to have an App-ID signature created.

### Identify Unknown Applications in Traffic Logs and View Packet Captures

Step 1 Verify that unknown application packet capture is enabled. This option is on by default.

1. To view the unknown application capture setting, run the following CLI command:

   ```
   admin@PA-200> show running application setting | match "Unknown capture"
   ```

2. If the unknown capture setting option is off, enable it:

   ```
   admin@PA-200> set application dump-unknown yes
   ```

Step 2 Locate unknown application by filtering the traffic logs.

1. Select **Monitor > Logs > Traffic**.
2. Click **Add Filter** and select the filters as shown in the following example.

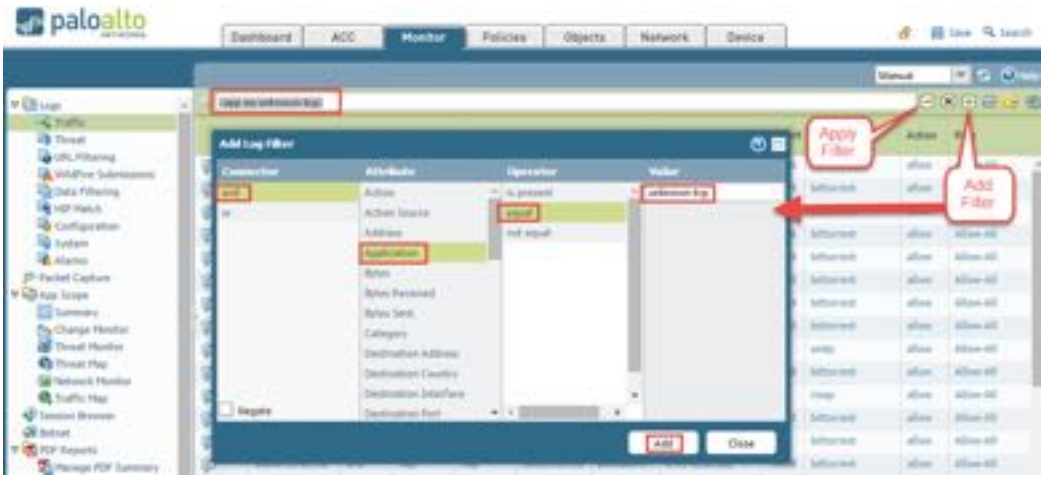


3. Click **Add** and **Apply Filter**.

Step 3 Click the packet capture icon to view the packet capture or **Export** it to your local system.

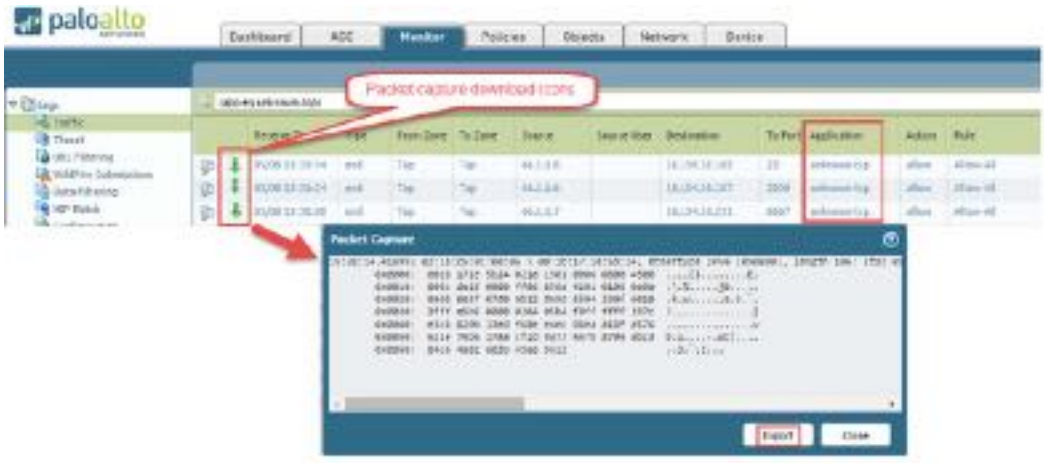

## Take a Custom Application Packet Capture

You can configure a Palo Alto Networks firewall to take a packet capture based on an application name and filters that you define. You can then use the packet capture to troubleshoot issues with controlling an application. When configuring an application packet capture, you must use the application name defined in the App-ID database. You can view a list of all App-ID applications using Applipedia or from the web interface on the firewall in **Objects > Applications**.

| Take a Custom Application Packet Capture | |
|---|---|
| Step 1 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall. |
| Step 2 | Turn on the application packet capture and define filters.<br>`admin@PA-200> set application dump on application <application-name> rule <rule-name>`<br>For example, to capture packets for the facebook-base application that matches the security rule named rule1, run the following CLI command:<br>`admin@PA-200> set application dump on application facebook-base rule rule1`<br> You can also apply other filters, such as source IP address and destination IP address. |

 © Palo Alto Networks, Inc.

**Take a Custom Application Packet Capture (Continued)**

Step 3 View the output of the packet capture settings to ensure that the correct filters are applied. The output appears after enabling the packet capture.

In the following output, you see that application filtering is now on based on the facebook-base application for traffic that matches rule1.

```
Application setting:
Application cache             : yes
Supernode                     : yes
Heuristics                    : yes
Cache Threshold               : 16
Bypass when exceeds queue limit: no
Traceroute appid              : yes
Traceroute TTL threshold      : 30
Use cache for appid           : no
Unknown capture               : on
Max. unknown sessions         : 5000
Current unknown sessions      : 0
Application capture           : on
Max. application sessions     : 5000
Current application sessions  : 0
Application filter setting:
    Rule                      : rule1
    From                      : any
    To                        : any
    Source                    : any
    Destination               : any
    Protocol                  : any
    Source Port               : any
    Dest. Port                : any
    Application               : facebook-base
Current APPID Signature
   Signature Usage            : 21  MB (Max. 32  MB)
      TCP 1 C2S               : 15503  states
      TCP 1 S2C               : 5070   states
      TCP 2 C2S               : 2426   states
      TCP 2 S2C               : 702    states
      UDP 1 C2S               : 11379  states
      UDP 1 S2C               : 2967   states
      UDP 2 C2S               : 755    states
      UDP 2 S2C               : 224    states
```

Step 4 To turn off application packet capture after the traffic you are interested in has traversed the firewall:

```
admin@PA-200> set application dump off
```

**Take a Custom Application Packet Capture (Continued)**

Step 5 View/export the packet capture.

1. Log in to the web interface on the firewall and select **Monitor > Logs > Traffic**.
2. In the log entry that you are interested in, click the green packet capture icon in the second column.

3. View the packet capture directly or **Export** it to your local system. The following screen capture shows the facebook-base packet capture.

## Take a Packet Capture on the Management Interface

The `tcpdump` CLI command enables you to capture packets that traverse the management interface (MGT) on a Palo Alto Networks firewall.

Each platform has a default number of bytes that `tcpdump` captures. The PA-200, PA-500, and PA-2000 Series firewalls capture 68 bytes of data from each packet and anything over that is truncated. The PA-3000, PA-4000, PA-5000 Series, the PA-7000 Series firewalls, and VM-Series firewalls capture 96 bytes of data from each packet. To define the number of packets that `tcpdump` will capture, use the `snaplen` (snap length) option (range 0-65535). Setting the `snaplen` to 0 will cause the firewall to use the maximum length required to capture whole packets.

**Take a Management Interface Packet Capture**

Step 1 Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall.

 © Palo Alto Networks, Inc.

**Take a Management Interface Packet Capture (Continued)**

Step 2 To start a packet capture on the MGT interface, run the following command:

```
admin@PA-200> tcpdump filter "<filter-option> <IP-address>" snaplen length
```

For example, to capture the traffic that is generated when and administrator authenticates to the firewall using RADIUS, filter on the destination IP address of the RADIUS server (10.5.104.99 in this example):

```
admin@PA-200> tcpdump filter "dst 10.5.104.99" snaplen 0
```

You can also filter on src (source IP address), host, net, and you can exclude content. For example, to filter on a subnet and exclude all SCP, SFTP, and SSH traffic (which uses port 22), run the following command:

```
admin@PA-200> tcpdump filter "net 10.5.104.0/24 and not port 22" snaplen 0
```

Each time `tcpdump` takes a packet capture, it stores the content in a file named mgmt.pcap. This file is overwritten each time you run `tcpdump`.

Step 3 After the traffic you are interested in has traversed the MGT interface, press `Ctrl` + `C` to stop the capture.

Step 4 View the packet capture by running the following command:

```
admin@PA-200> view-pcap mgmt-pcap mgmt.pcap
```

The following output shows the packet capture from the MGT port (10.5.104.98) to the RADIUS server (10.5.104.99):

```
09:55:29.139394 IP 10.5.104.98.43063 > 10.5.104.99.radius: RADIUS, Access Request (1), id: 0x00 length: 89
09:55:29.144354 arp reply 10.5.104.98 is-at 00:25:90:23:94:98 (oui Unknown)
09:55:29.379290 IP 10.5.104.98.43063 > 10.5.104.99.radius: RADIUS, Access Request (1), id: 0x00 length: 70
09:55:34.379262 arp who-has 10.5.104.99 tell 10.5.104.98
```

Step 5 (Optional) Export the packet capture from the firewall using SCP (or TFTP). For example, to export the packet capture using SCP, run the following command:

```
admin@PA-200> scp export mgmt-pcap from mgmt.pcap to username@host:path
```

For example, to export the pcap to an SCP enabled server at 10.5.5.20 to a temp folder named temp-SCP, run the following CLI command:

```
admin@PA-200> scp export mgmt-pcap from mgmt.pcap to admin@10.5.5.20:c:/temp-SCP
```

Enter the login name and password for the account on the SCP server and the firewall will copy the packet capture to the SCP enabled server to c:\temp-SCP.

Step 6 You can now view the packet capture files using a network packet analyzer, such as Wireshark.

# Monitor Applications and Threats

All Palo Alto Networks next-generation firewalls come equipped with the App-ID technology, which identifies the applications traversing your network, irrespective of protocol, encryption, or evasive tactic. You can then Use the Application Command Center to monitor the applications. The ACC graphically summarizes the data from a variety of log databases to highlight the applications traversing your network, who is using them, and their potential security impact. ACC is dynamically updated, using the continuous traffic classification that App-ID performs; if an application changes ports or behavior, App-ID continues to see the traffic, displaying the results in ACC. Additional visibility into URL categories, threats, and data provides a complete and well-rounded picture of network activity. With ACC, you can very quickly learn more about the traffic traversing the network and then translate that information into a more informed security policy

You can also Use the Dashboard to monitor the network.

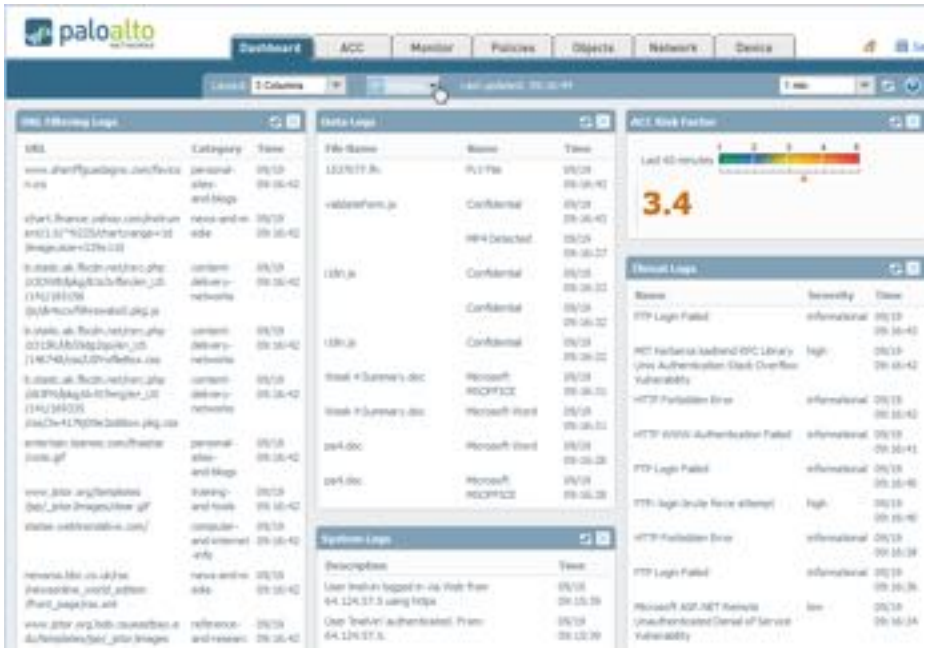


 © Palo Alto Networks, Inc.

# Monitor and Manage Logs

All Palo Alto Networks next-generation firewalls can generate log files that provide an audit trail of the activities and events on the firewall. There are separate logs for separate types of activities and events. For example, the Threat logs record all traffic that causes the firewall to generate a security alarm, URL Filtering logs record all traffic that matches a URL Filtering profile attached to a security rule, and Config logs record all changes to the firewall configuration. The firewall uses log data to generate reports (**Monitor > Reports**), which display the data in a tabular or graphical format. See Manage Reporting for details.

The following topics describe how to view logs locally on the firewall. You can also Use External Services for Monitoring logs.

- View the Log Files
- Filter Log Data
- Configure Log Storage Quotas and Expiration Periods
- Log Severity Levels and WildFire Verdicts
- Schedule Log Exports to an SCP or FTP Server

© Palo Alto Networks, Inc.

## View the Log Files

The firewall maintains logs for WildFire, configurations, system events, alarms, traffic flows, threats, URL filtering, data filtering, and Host Information Profile (HIP) matches. You can view the current logs at any time. To locate specific entries, you can apply filters to most of the log fields.

The firewall displays the information in logs so that role-based administration permissions are respected. When you display logs, only the information that you have permission to see is included. For information on administrator permissions, see Administrative Roles.

By default all log files are generated and stored locally on the firewall. You can view these log files in the **Monitor > Logs** pages:

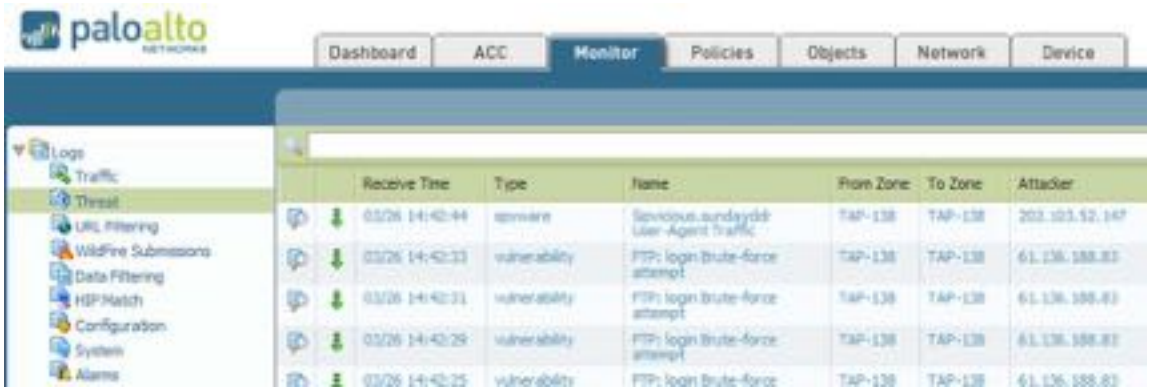


To display additional details, click the spyglass icon for an entry.

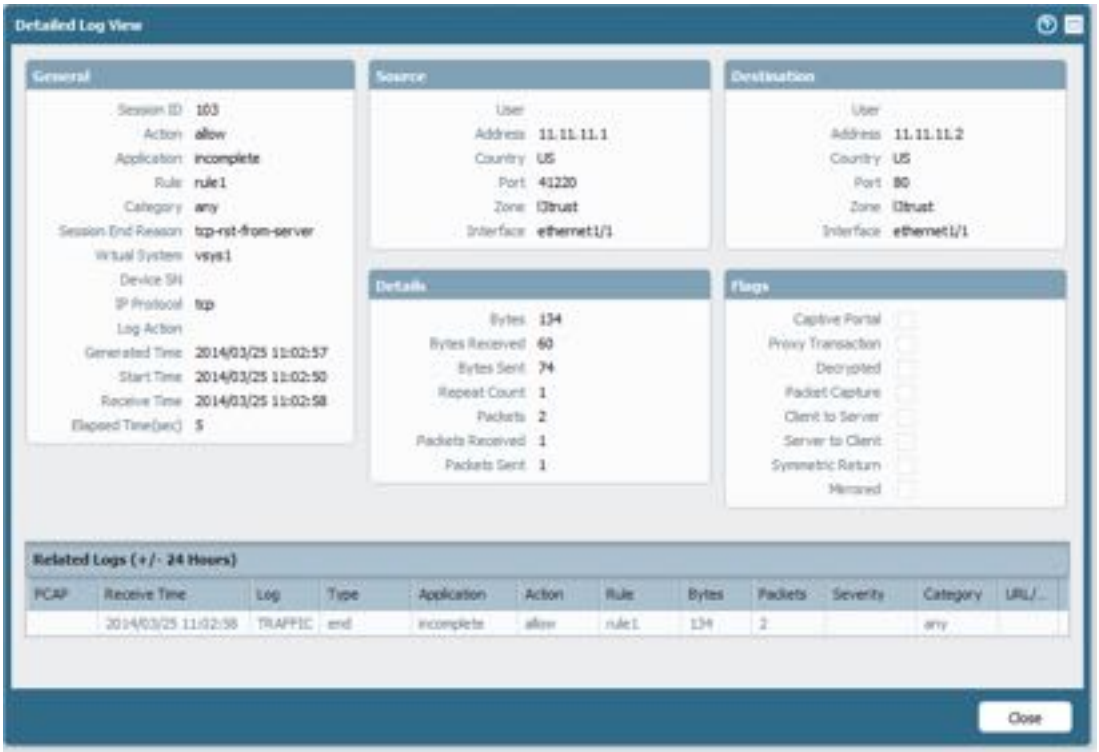


 © Palo Alto Networks, Inc.

The following table includes information on each log type:

| Log Description Charts | Description |
|---|---|
| Traffic | Displays an entry for the start and end of each session. Each entry includes the date and time, source and destination zones, addresses and ports, application name, security rule name applied to the flow, rule action (**allow**, **deny**, or **drop**), ingress and egress interface, number of bytes, and session end reason.<br>Click next to an entry to view additional details about the session, such as whether an ICMP entry aggregates multiple sessions between the same source and destination (the **Count** value will be greater than one).<br>The **Type** column indicates whether the entry is for the start or end of the session, or whether the session was denied or dropped. A **drop** indicates that the security rule that blocked the traffic specified **any** application, while a **deny** indicates the rule identified a specific application.<br>If traffic is dropped before the application is identified, such as when a rule drops all traffic for a specific service, the application is shown as **not-applicable**. |
| Threat | Displays an entry when traffic matches a Security Profile (Antivirus, Anti-Spyware, Vulnerability, URL Filtering, File Blocking, Data Filtering, or DoS Protection) that is attached to a security rule on the firewall. Each entry includes the date and time, a threat name or URL, the source and destination zones, addresses, and ports, the application name, and the alarm action (**allow** or **block**) and severity.<br>Click next to an entry to view additional details about the threat, such as whether the entry aggregates multiple threats of the same type between the same source and destination (the **Count** value will be greater than one).<br>The **Type** column indicates the type of threat, such as "virus" or "spyware." The Name column is the threat description or URL, and the **Category** column is the threat category (such as "keylogger") or URL category.<br>If local packet captures are enabled, click next to an entry to access the captured packets. To enable local packet captures, see Take Packet Captures.<br>For details on threat severity levels, see Log Severity Levels and WildFire Verdicts. |
| URL Filtering | Displays logs for all traffic that matches a URL Filtering profile attached to a security rule. For example, if rule blocks access to specific web sites and web site categories or if rule is configured to generate an alert when a web site is accessed. For information on defining URL filtering profiles, see URL Filtering. |
| WildFire Submissions | Displays logs for files that are uploaded and analyzed by the WildFire cloud; log data is sent back to the device after analysis, along with the analysis results.<br>For details on WildFire verdicts (benign or malicious), see Log Severity Levels and WildFire Verdicts. |

| Log Description Charts | Description |
|---|---|
| Data Filtering | Displays logs for the security rules that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall. See Set Up Data Filtering for information on defining data filtering profiles.<br>This log also shows information for file-blocking profiles. For example, if you are blocking .exe files, the log will show the files that were blocked. If you forward files to WildFire, you will see the results of that action. In this case, if you are forwarding PE files to WildFire, for example, the log will show that the file was forwarded and will also show the status on whether or not it was uploaded to WildFire successfully. |
| Configuration | Displays an entry for each configuration change. Each entry includes the date and time, the administrator username, the IP address from where the change was made, the type of client (XML, Web or CLI), the type of command executed, whether the command succeeded or failed, the configuration path, and the values before and after the change. |
| System | Displays an entry for each system event. Each entry includes the date and time, the event severity, and an event description.<br>For details on System log severity levels, see Log Severity Levels and WildFire Verdicts. |
| HIP Match | Displays traffic flows that match a HIP Object or HIP Profile that you have configured. |

 © Palo Alto Networks, Inc.

## Filter Log Data

Each log page has a filter area at the top of the page.

Use the filter area as follows:

- Click any of the underlined links in the log listing to add that item as a log filter option. For example, if you click the **Host** link in the log entry for 10.0.0.252 and **Web Browsing**, both items are added, and the search will find entries that match both (AND search).
- To define other search criteria, click **Add Log Filter**. Select the type of search (and/or), the attribute to include in the search, the matching operator, and the values for the match, if appropriate. Click **Add** to add the criterion to the filter area on the log page, and then click **Close** to close the pop-up window. Click **Apply Filter** to display the filtered list.

If the **Value** string matches an **Operator** (such as **has** or **in**), enclose the string in quotation marks to avoid a syntax error. For example, if you filter by destination country and use `IN` as a **Value** to specify INDIA, enter the filter as `( dstloc eq "IN" )`.

You can combine filter expressions added on the log page with those you define in the Add Log Filter dialog. The filter field on the log page displays each filter as an entry.

If you add a **Receive Time** filter with the **Operator** set to **in** and the **Value** set to **Last 60 seconds**, some of the page links on the log viewer might not show results because the number of pages might grow or shrink due to the dynamic nature of the selected time.

- To clear filters and redisplay the unfiltered list, click **Clear Filter**.
- To save your selections as a new filter, click **Save Filter**, enter a name for the filter, and click **OK**.
- To export the current log listing (as shown on the page, including any applied filters) click **Save Filter**. Select whether to open the file or save it to disk, and select the check box if you want to always use the same option. Click **OK**.
- To export the current log listing in CSV Format, select the **Export to CSV** icon. By default, exporting the log listing to CSV format generates a CSV report with up to 2,000 rows of logs. To change the limit for rows displayed in CSV reports, use the **Max Rows in CSV Export** field on the **Log Export and Reporting** tab (select **Device > Setup > Management > Logging and Reporting Settings**).

To change the automatic refresh interval, select an interval from the drop-down (**1 min**, **30 seconds**, **10 seconds**, or **Manual**).

To change the number of log entries per page, select the number of rows from the **Rows** drop-down.

Log entries are retrieved in blocks of 10 pages. Use the paging controls at the bottom of the page to navigate through the log list. Select the **Resolve Hostname** check box to begin resolving external IP addresses to domain names.

© Palo Alto Networks, Inc.

## Configure Log Storage Quotas and Expiration Periods

The firewall automatically deletes logs that exceed the expiration period. When the firewall reaches the storage quota for a log type, it automatically deletes older logs of that type to create space even if you don't set an expiration period.

If you want to manually delete logs, select **Device > Log Settings** and, in the Manage Logs section, click the links to clear logs by type.

| Configure Log Storage Quotas and Expiration Periods | |
|---|---|
| Step 1 | Select **Device > Setup > Management** and edit the Logging and Reporting Settings. |
| Step 2 | In the **Log Storage** tab, select the **Log Storage** check box and enter the **Quota (%)** for each log type. When you change a percentage value, the dialog refreshes to display the corresponding absolute value (Quota GB/MB column). |
| Step 3 | Enter the **Max Days** (expiration period) for each log type (range is 1-2,000). The fields are blank by default, which means the logs never expire.<br> The firewall synchronizes expiration periods across high availability (HA) pairs. Because only the active HA peer generates logs, the passive peer has no logs to delete unless failover occurs and it starts generating logs. |
| Step 4 | Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

## Log Severity Levels and WildFire Verdicts

The following table summarizes the Threat severity levels:

| Severity | Description |
|---|---|
| **Critical** | Serious threats, such as those that affect default installations of widely deployed software, result in root compromise of servers, and the exploit code is widely available to attackers. The attacker usually does not need any special authentication credentials or knowledge about the individual victims and the target does not need to be manipulated into performing any special functions. |
| **High** | Threats that have the ability to become critical but have mitigating factors; for example, they may be difficult to exploit, do not result in elevated privileges, or do not have a large victim pool. |
| **Medium** | Minor threats in which impact is minimized, such as DoS attacks that do not compromise the target or exploits that require an attacker to reside on the same LAN as the victim, affect only non-standard configurations or obscure applications, or provide very limited access. In addition, WildFire Submissions log entries with a malware verdict are logged as Medium. |
| **Low** | Warning-level threats that have very little impact on an organization's infrastructure. They usually require local or physical system access and may often result in victim privacy or DoS issues and information leakage. Data Filtering profile matches are logged as Low. |
| **Informational** | Suspicious events that do not pose an immediate threat, but that are reported to call attention to deeper problems that could possibly exist. URL Filtering log entries and WildFire Submissions log entries with a benign verdict are logged as Informational. |

The following table summarizes the System log severity levels. For a partial list of System log messages and their corresponding severity levels, refer to System Log Events.

| Severity | Description |
|---|---|
| **Critical** | Hardware failures, including HA failover and link failures. |
| **High** | Serious issues, including dropped connections with external devices, such as LDAP and RADIUS servers. |
| **Medium** | Mid-level notifications, such as antivirus package upgrades. |
| **Low** | Minor severity notifications, such as user password changes. |
| **Informational** | Log in/log off, administrator name or password change, any configuration change, and all other events not covered by the other severity levels. |

The following table summarizes the Correlation log severity levels:

| Severity | Description |
|---|---|
| **Critical** | Confirms that a host has been compromised based on correlated events that indicate an escalation pattern. For example, a critical event is logged when a host that received a file with a malicious verdict by WildFire, exhibits the same command-and control activity that was observed in the WildFire sandbox for that malicious file. |

| Severity | Description |
|---|---|
| **High** | Indicates that a host is very likely compromised based on a correlation between multiple threat events, such as malware detected anywhere on the network that matches the command and control activity being generated from a particular host. |
| **Medium** | Indicates that a host is likely compromised based on the detection of one or multiple suspicious events, such as repeated visits to known malicious URLs that suggests a scripted command-and-control activity. |
| **Low** | Indicates that a host is possibly compromised based on the detection of one or multiple suspicious events, such as a visit to a malicious URL or a dynamic DNS domain. |
| **Informational** | Detects an event that may be useful in aggregate for identifying suspicious activity; each event is not necessarily significant on its own. |

The following table summarizes the WildFire verdicts:

| Severity | Description |
|---|---|
| **Benign** | Indicates that the entry received a WildFire analysis verdict of benign. Files categorized as benign are safe and do not exhibit malicious behavior. |
| **Grayware** | Indicates that the entry received a WildFire analysis verdict of grayware. Files categorized as grayware do not pose a direct security threat, but might display otherwise obtrusive behavior. Grayware can include, adware, spyware, and Browser Helper Objects (BHOs). |
| **Malware** | Indicates that the entry received a WildFire analysis verdict of malware. Files categorized as malware are malicious in intent or nature and can pose a security threat. Malware can include viruses, worms, Trojans, Remote Access Tools (RATs), rootkits, and botnets. For files that are identified as malware, a signature is generated and distributed by the WildFire cloud to prevent against future exposure. |

 © Palo Alto Networks, Inc.

# Schedule Log Exports to an SCP or FTP Server

You can schedule exports of Traffic, Threat, URL Filtering, Data Filtering, HIP Match, and WildFire Submission logs to a Secure Copy (SCP) server or File Transfer Protocol (FTP) server. Perform this task for each log type you want to export.

You can use Secure Copy (SCP) commands from the CLI to export the entire log database to an SCP server and import it to another firewall. Because the log database is too large for an export or import to be practical on the following platforms, they do not support these options: PA-7000 Series firewalls (all PAN-OS releases), Panorama virtual appliance running Panorama 6.0 or later releases, and Panorama M-Series appliances (all Panorama releases).

| Schedule Log Exports to an SCP or FTP Server | |
|---|---|
| Step 1 | Select **Device > Scheduled Log Export** and click **Add**. |
| Step 2 | Enter a **Name** for the scheduled log export and **Enable** it. |
| Step 3 | Select the **Log Type** to export. |
| Step 4 | Select the daily **Scheduled Export Start Time**. The options are in 15-minute increments for a 24-hour clock (00:00 - 23:59). |
| Step 5 | Select the **Protocol** to export the logs: **SCP** (secure) or **FTP**. |
| Step 6 | Enter the **Hostname** or IP address of the server. |
| Step 7 | Enter the **Port** number. By default, FTP uses port 21 and SCP uses port 22. |
| Step 8 | Enter the **Path** or directory in which to save the exported logs. |
| Step 9 | Enter the **Username** and, if necessary, the **Password** (and **Confirm Password**) to access the server. |
| Step 10 | (FTP only) Select the **Enable FTP Passive Mode** check box if you want to use FTP passive mode, in which the firewall initiates a data connection with the FTP server. By default, the firewall uses FTP active mode, in which the FTP server initiates a data connection with the firewall. Choose the mode based on what your FTP server supports and on your network requirements. |
| Step 11 | (SCP only) Click **Test SCP server connection**. The connection is not established until the firewall accepts the host key for the SCP server. |
| Step 12 | Click **OK** and **Commit**. |

# Manage Reporting

The reporting capabilities on the firewall allow you to keep a pulse on your network, validate your policies, and focus your efforts on maintaining network security for keeping your users safe and productive.


- Report Types
- View Reports
- Configure the Report Expiration Period
- Disable Predefined Reports
- Generate Custom Reports
- Generate Botnet Reports
- Manage PDF Summary Reports
- Generate User/Group Activity Reports
- Manage Report Groups
- Schedule Reports for Email Delivery


 © Palo Alto Networks, Inc.

## Report Types

The firewall includes predefined reports that you can use as-is, or you can build custom reports that meet your needs for specific data and actionable tasks, or you can combine predefined and custom reports to compile information you need. The firewall provides the following types of reports:

- **Predefined Reports—**Allow you to view a quick summary of the traffic on your network. A suite of predefined reports are available in four categories—Applications, Traffic, Threat, and URL Filtering. See View Reports.
- **User or Group Activity Reports—**Allow you to schedule or create an on-demand report on the application use and URL activity for a specific user or for a user group. The report includes the URL categories and an estimated browse time calculation for individual users. See Generate User/Group Activity Reports.
- **Custom Reports**—Create and schedule custom reports that show exactly the information you want to see by filtering on conditions and columns to include. You can also include query builders for more specific drill down on report data. See Generate Custom Reports.
- **PDF Summary Reports**—Aggregate up to 18 predefined or custom reports/graphs from Threat, Application, Trend, Traffic, and URL Filtering categories into one PDF document. See Manage PDF Summary Reports.
- **Botnet Reports**—Allow you to use behavior-based mechanisms to identify potential botnet-infected hosts in the network. See Generate Botnet Reports.
- **Report Groups**—Combine custom and predefined reports into report groups and compile a single PDF that is emailed to one or more recipients. See Manage Report Groups.

Reports can be generated on demand, on a recurring schedule, and can be scheduled for email delivery.

# View Reports

The firewall provides an assortment of over 40 predefined reports that it generates every day. You can view these reports directly on the firewall. You can also view custom reports and summary reports.

About 200 MB of storage is allocated for saving reports on the firewall. You can't configure this limit but you can Configure the Report Expiration Period: the firewall will automatically delete reports that exceed the period. Keep in mind that when the firewall reaches its storage limit, it automatically deletes older reports to create space even if you don't set an expiration period. Another way to conserve system resources on the firewall is to Disable Predefined Reports. For long-term retention of reports, you can export the reports (as described below) or Schedule Reports for Email Delivery.

Unlike other reports, you can't save User/Group Activity reports on the firewall. You must Generate User/Group Activity Reports on demand or schedule them for email delivery.

**View Reports**

Step 1 Select **Monitor > Reports**.

The reports are chunked into sections on the right-hand side of the window: **Custom Reports**, **Application Reports**, **Traffic Reports**, **Threat Reports**, **URL Filtering Reports**, and **PDF Summary Reports**.

Step 2 Select a report to view. When you select a report, the previous day's report is displayed onscreen.

To view reports for any of the previous days, select an available date from the calendar at the bottom of the page and select a report within the same section. If you change sections, the time selection is reset.

Step 3 To view a report offline, you can export the report to PDF, CSV or to XML formats. Click **Export to PDF**, **Export to CSV**, or **Export to XML** at the bottom of the page. Then print or save the file.

 © Palo Alto Networks, Inc.

## Configure the Report Expiration Period

When you set the **Report Expiration Period**, it applies to all Report Types. The firewall automatically deletes reports that exceed the period.

| Configure Report Expiration Periods | |
|---|---|
| Step 1 | Select **Device > Setup > Management**, edit the Logging and Reporting Settings, and select the **Log Export and Reporting** tab. |
| Step 2 | Enter the **Report Expiration Period** in days (range is 1-2000, default is no expiration).<br> You can't change the storage that the firewall allocates for saving reports: it is predefined at about 200 MB. When the firewall reaches the storage maximum, it automatically deletes older reports to create space even if you don't set a **Report Expiration Period**. |
| Step 3 | Click **OK** and **Commit**. |

# Disable Predefined Reports

The firewall includes about 40 predefined reports that it automatically generates daily. If you do not use some or all of these, you can disable selected reports to conserve system resources on the firewall.

Make sure that no report group or PDF summary report includes the predefined reports you will disable. Otherwise, the firewall will render the PDF summary report or report group without any data.

**Disable Predefined Reports**

Step 1 Select **Device > Setup > Management** and edit the Logging and Reporting Settings.

Step 2 Select the **Pre-Defined Reports** tab and clear the check box for each report you want to disable. To disable all predefined reports, click **Deselect All**.

Step 3 Click **OK** and **Commit**.

 © Palo Alto Networks, Inc.

## Generate Custom Reports

In order to create purposeful custom reports, you must consider the attributes or key pieces of information that you want to retrieve and analyze. This consideration guides you in making the following selections in a custom report:

| Selection | Description |
|---|---|
| Data Source | The data file that is used to generate the report. The firewall offers two types of data sources—Summary databases and Detailed logs.<br>• Summary databases are available for traffic, threat, and application statistics. The firewall aggregates the detailed logs on traffic, application, and threat at 15-minute intervals. The data is condensed—duplicate sessions are grouped together and incremented with a repeat counter, and some attributes (or columns) are not included in the summary—to allow faster response time when generating reports.<br>• Detailed logs are itemized and are a complete listing of all the attributes (or columns) that pertain to the log entry. Reports based on detailed logs take much longer to run and are not recommended unless absolutely necessary. |
| Attributes | The columns that you want to use as the match criteria. The attributes are the columns that are available for selection in a report. From the list of **Available Columns**, you can add the selection criteria for matching data and for aggregating the details (the **Selected Columns**). |
| Sort By/ Group By | The **Sort By** and the **Group By** criteria allow you to organize/segment the data in the report; the sorting and grouping attributes available vary based on the selected data source.<br>The Sort By option specifies the attribute that is used for aggregation. If you do not select an attribute to sort by, the report will return the first N number of results without any aggregation.<br>The Group By option allows you to select an attribute and use it as an anchor for grouping data; all the data in the report is then presented in a set of top 5, 10, 25 or 50 groups. For example, when you select Hour as the Group By selection and want the top 25 groups for a 24-hr time period, the results of the report will be generated on an hourly basis over a 24-hr period. The first column in the report will be the hour and the next set of columns will be the rest of your selected report columns. |

| Selection | Description |
|---|---|
| | The following example illustrates how the **Selected Columns** and **Sort By**/**Group By** criteria work together when generating reports:<br> 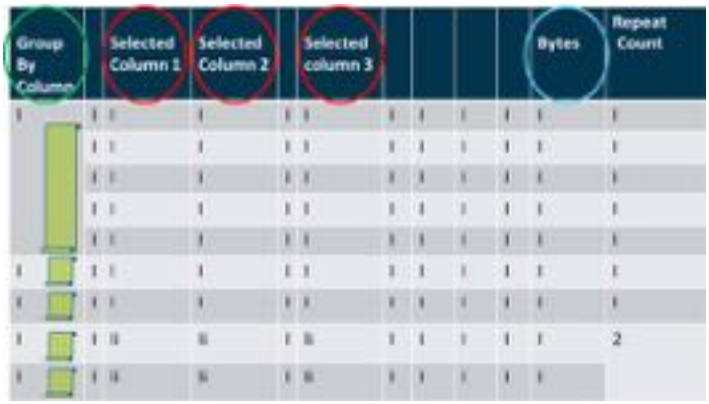 <br>The columns circled in red (above) depict the columns selected, which are the attributes that you match against for generating the report. Each log entry from the data source is parsed and these columns are matched on. If multiple sessions have the same values for the selected columns, the sessions are aggregated and the repeat count (or sessions) is incremented.<br>The column circled in blue indicates the chosen sort order. When the sort order (**Sort By**) is specified, the data is sorted (and aggregated) by the selected attribute.<br>The column circled in green indicates the **Group By** selection, which serves as an anchor for the report. The **Group By** column is used as a match criteria to filter for the top N groups. Then, for each of the top N groups, the report enumerates the values for all the other selected columns. |

 © Palo Alto Networks, Inc.

| Selection | Description |
|---|---|
| | For example, if a report has the following selections<br> 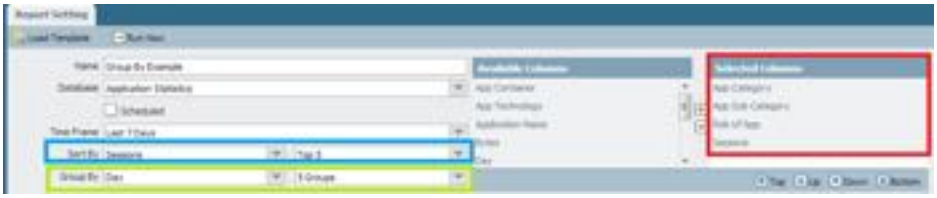 <br>The output will display as follows:<br> <br>The report is anchored by **Day** and sorted by **Sessions**. It lists the 5 days (**5 Groups**) with maximum traffic in the **Last 7 Days** time frame. The data is enumerated by the **Top 5** sessions for each day for the selected columns—**App Category**, **App Subcategory** and **Risk**. |
| Time Period | The date range for which you want to analyze data. You can define a custom range or select a time period ranging from last 15 minutes to the last 30 days. The reports can be run on demand or scheduled to run at a daily or weekly cadence. |
| Query Builder | The query builder allows you to define specific queries to further refine the selected attributes. It allows you see just what you want in your report using **and** and **or** operators and a match criteria, and then include or exclude data that matches or negates the query in the report. Queries enable you to generate a more focused collation of information in a report. |

| Generate Custom Reports | |
|---|---|
| Step 1 | Select **Monitor > Manage Custom Reports**. |
| Step 2 | Click **Add** and then enter a **Name** for the report.<br>To base a report on an predefined template, click **Load Template** and choose the template. You can then edit the template and save it as a custom report. |

**Generate Custom Reports**

Step 3 Select the **Database** to use for the report.

Each time you create a custom report, a **Log View** report is automatically created. This report show the logs that were used to build the custom report. The log view report uses the same name as the custom report, but appends the phrase (Log View) to the report name.

When creating a report group, you can include the log view report with the custom report. For more information, see Manage Report Groups.

Step 4 Select the **Scheduled** check box to run the report each night. The report is then available for viewing in the **Reports** column on the side.

Step 5 Define the filtering criteria. Select the **Time Frame**, the **Sort By** order, **Group By** preference, and select the columns that must display in the report.

Step 6 (Optional) Select the **Query Builder** attributes if you want to further refine the selection criteria. To build a report query, specify the following and click **Add**. Repeat as needed to construct the full query.

- **Connector**—Choose the connector (and/or) to precede the expression you are adding.
- **Negate**—Select the check box to interpret the query as a negation. If, for example, you choose to match entries in the last 24 hours and/or are originating from the untrust zone, the negate option causes a match on entries that are not in the past 24 hours and/or are not from the untrust zone.
- **Attribute**—Choose a data element. The available options depend on the choice of database.
- **Operator**—Choose the criterion to determine whether the attribute applies (such as =). The available options depend on the choice of database.
- **Value**—Specify the attribute value to match.

For example, the following figure (based on the Traffic Log database) shows a query that matches if the traffic log entry was received in the past 24 hours and is from the "untrust" zone.

Step 7 To test the report settings, select **Run Now**. Modify the settings as required to change the information that is displayed in the report.

Step 8 Click **OK** to save the custom report.

 © Palo Alto Networks, Inc.

## Generate Custom Reports

### Examples of Custom Reports

If you want to set up a simple report in which you use the traffic summary database from the last 30 days, and sort the data by the top 10 sessions and these sessions are grouped into 5 groups by day of the week. You would set up the custom report to look like this:

And the PDF output for the report would look as follows:

## Generate Custom Reports

Now, if you want to use the query builder to generate a custom report that represents the top consumers of network resources within a user group, you would set up the report to look like this:

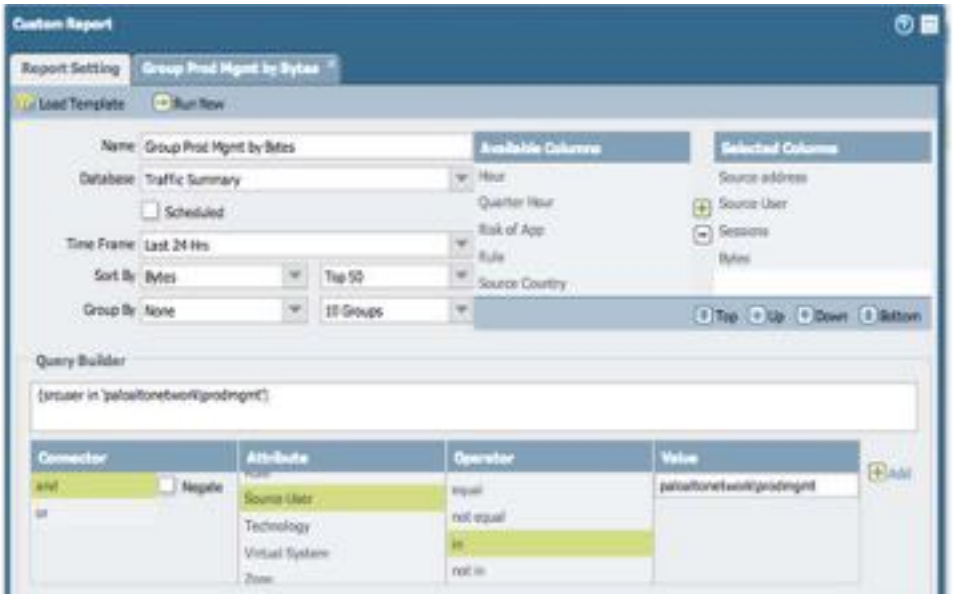


The report would display the top users in the product management user group sorted by bytes, as follows:

 © Palo Alto Networks, Inc.

# Generate Botnet Reports

The botnet report enables you to use heuristic and behavior-based mechanisms to identify potential malware- or botnet-infected hosts in your network. To evaluate botnet activity and infected hosts, the firewall correlates user and network activity data in Threat, URL, and Data Filtering logs with the list of malware URLs in PAN-DB, known dynamic DNS domain providers, and domains registered within the last 30 days. You can configure the report to identify hosts that visited those sites, as well as hosts that communicated with Internet Relay Chat (IRC) servers or that used unknown applications. Malware often use dynamic DNS to avoid IP blacklisting, while IRC servers often use bots for automated functions.

The firewall requires Threat Prevention and URL Filtering licenses to use the botnet report.
You can Use the Automated Correlation Engine to monitor suspicious activities based on additional indicators besides those that the botnet report uses. However, the botnet report is the only tool that uses newly registered domains as an indicator.

- ▲ Configure a Botnet Report
- ▲ Interpret Botnet Report Output

## Configure a Botnet Report

You can schedule a botnet report or run it on demand. The firewall generates scheduled botnet reports every 24 hours because behavior-based detection requires correlating traffic across multiple logs over that timeframe.

| **Configure a Botnet Report** | | |
|---|---|---|
| Step 1 | Define the types of traffic that indicate possible botnet activity. | 1. Select **Monitor > Botnet** and click **Configuration** on the right side of the page.<br>2. **Enable** and define the **Count** for each type of HTTP Traffic that the report will include.<br>The **Count** values represent the minimum number of events of each traffic type that must occur for the report to list the associated host with a higher confidence score (higher likelihood of botnet infection). If the number of events is less than the **Count**, the report will display a lower confidence score or (for certain traffic types) won't display an entry for the host. For example, if you set the **Count** to three for **Malware URL visit**, then hosts that visit three or more known malware URLs will have higher scores than hosts that visit less than three. For details, see Interpret Botnet Report Output.<br>3. Define the thresholds that determine whether the report will include hosts associated with traffic involving Unknown TCP or Unknown UDP applications.<br>4. Select the **IRC** check box to include traffic involving IRC servers.<br>5. Click **OK** to save the report configuration. |

| Configure a Botnet Report (Continued) | |
|---|---|
| Step 2 Schedule the report or run it on demand. | 1. Click **Report Setting** on the right side of the page.<br>2. Select a time interval for the report in the **Test Run Time Frame** drop-down.<br>3. Select the **No. of Rows** to include in the report.<br>4. (Optional) **Add** queries to the Query Builder to filter the report output by attributes such as source/destination IP addresses, users, or zones.<br>For example, if you know in advance that traffic initiated from the IP address 10.3.3.15 contains no potential botnet activity, you can add `not (addr.src in 10.0.1.35)` as a query to exclude that host from the report output. For details, see Interpret Botnet Report Output.<br>5. Select **Scheduled** to run the report daily or click **Run Now** to run the report immediately.<br>6. Click **OK** and **Commit**. |

## Interpret Botnet Report Output

The botnet report displays a line for each host that is associated with traffic you defined as suspicious when configuring the report. For each host, the report displays a confidence score of 1 to 5 to indicate the likelihood of botnet infection, where 5 indicates the highest likelihood. The scores correspond to threat severity levels: 1 is informational, 2 is low, 3 is medium, 4 is high, and 5 is critical. The firewall bases the scores on:

- **Traffic type**—Certain HTTP traffic types are more likely to involve botnet activity. For example, the report assigns a higher confidence to hosts that visit known malware URLs than to hosts that browse to IP domains instead of URLs, assuming you defined both those activities as suspicious.
- **Number of events**—Hosts that are associated with a higher number of suspicious events will have higher confidence scores based on the thresholds (**Count** values) you define when you Configure a Botnet Report.
- **Executable downloads**—The report assigns a higher confidence to hosts that download executable files. Executable files are a part of many infections and, when combined with the other types of suspicious traffic, can help you prioritize your investigations of compromised hosts.

When reviewing the report output, you might find that the sources the firewall uses to evaluate botnet activity (for example, the list of malware URLs in PAN-DB) have gaps. You might also find that these sources identify traffic that you consider safe. To compensate in both cases, you can add query filters when you Configure a Botnet Report.

 © Palo Alto Networks, Inc.

# Manage PDF Summary Reports

PDF summary reports contain information compiled from existing reports, based on data for the top 5 in each category (instead of top 50). They also contain trend charts that are not available in other reports.

| Generate PDF Summary Reports | |
|---|---|
| Step 1 Set up a **PDF Summary Report**. | 1. Select **Monitor > PDF Reports > Manage PDF Summary**.<br>2. Click **Add** and then enter a **Name** for the report.<br>3. Use the drop-down for each report group and select one or more of the elements to design the PDF Summary Report. You can include a maximum of 18 report elements.<br> 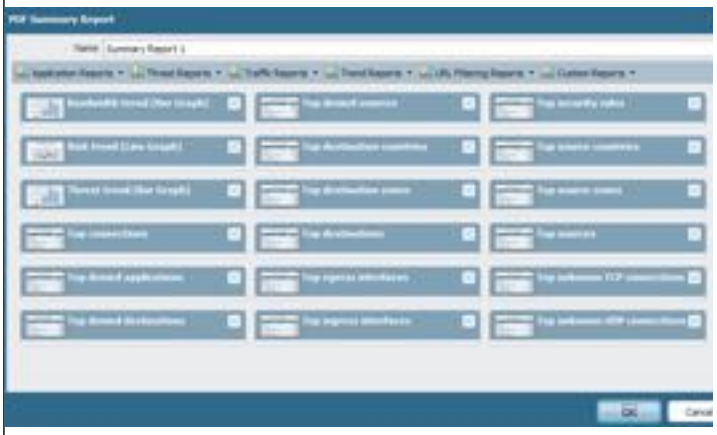 <br>• To remove an element from the report, click the **x** icon or clear the selection from the drop-down for the appropriate report group.<br>• To rearrange the reports, drag and drop the icons to another area of the report.<br>4. Click **OK t**o save the report.<br>5. **Commit** the changes. |

| Generate PDF Summary Reports | |
|---|---|
| Step 2 View the report. | To download and view the PDF Summary Report, see View Reports. 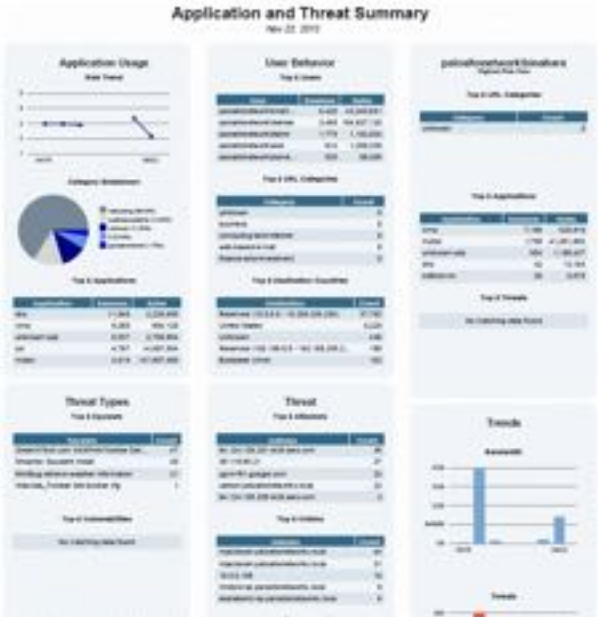  |

 © Palo Alto Networks, Inc.

# Generate User/Group Activity Reports

User/Group Activity reports summarize the web activity of individual users or user groups. Both reports include the same information except for the **Browsing Summary by URL Category** and **Browse time calculations**, which only the User Activity report includes.

You must configure User-ID on the firewall to access the list of users and user groups.

| **Generate User/Group Activity Reports** | | |
|---|---|---|
| Step 1 | Configure the browse times and number of logs for User/Group Activity reports.<br><br>Required only if you want to change the default values. | 1. Select **Device > Setup > Management**, edit the Logging and Reporting Settings, and select the **Log Export and Reporting** tab.<br>2. For the **Max Rows in User Activity Report**, enter the maximum number of rows that the detailed user activity report supports (range is 1-1048576, default is 5000). This determines the number of logs that the report analyzes.<br>3. Enter the **Average Browse Time** in seconds that you estimate users should take to browse a web page (range is 0-300, default is 60). Any request made after the average browse time elapses is considered a new browsing activity. The calculation uses Container Pages (logged in the URL Filtering logs) as the basis and ignores any new web pages that are loaded between the time of the first request (start time) and the average browse time. For example, if you set the **Average Browse Time** to two minutes and a user opens a web page and views that page for five minutes, the browse time for that page will still be two minutes. This is done because the firewall can't determine how long a user views a given page. The average browse time calculation ignores sites categorized as web advertisements and content delivery networks.<br>4. For the **Page Load Threshold**, enter the estimated time in seconds for page elements to load on the page (default is 20). Any requests that occur between the first page load and the page load threshold are assumed to be elements of the page. Any requests that occur outside of the page load threshold are assumed to be the user clicking a link within the page.<br>5. Click **OK** to save your changes. |

| Generate User/Group Activity Reports (Continued) | |
|---|---|
| Step 2 Generate the User/Group Activity report. | 1. Select **Monitor > PDF Reports > User Activity Report**.<br>2. Click **Add** and then enter a **Name** for the report.<br>3. Create the report:<br>• User Activity Report—Select **User** and enter the **Username** or **IP address** (IPv4 or IPv6) of the user.<br>• Group Activity Report—Select **Group** and select the **Group Name** of the user group.<br>4. Select the **Time Period** for the report.<br>5. Optionally, select the **Include Detailed Browsing** check box (default is cleared) to include detailed URL logs in the report.<br>The detailed browsing information can include a large volume of logs (thousands of logs) for the selected user or user group and can make the report very large.<br>6. To run the report on demand, click **Run Now**.<br>7. To save the report configuration, click **OK**. You can't save the output of User/Group Activity reports on the firewall. To schedule the report for email delivery, see Schedule Reports for Email Delivery. |

 © Palo Alto Networks, Inc.

# Manage Report Groups

Report groups allow you to create sets of reports that the system can compile and send as a single aggregate PDF report with an optional title page and all the constituent reports included.

| Set up Report Groups | |
|---|---|
| Step 1 Set up report groups.<br><br>You must set up a **Report Group** to email report(s). | 1. Create an Email server profile.<br>2. Define the **Report Group**. A report group can compile predefined reports, PDF Summary reports, custom reports, and Log View report into a single PDF.<br>a. Select **Monitor > Report Group**.<br>b. Click **Add** and then enter a **Name** for the report group.<br>c. (Optional) Select **Title Page** and add a **Title** for the PDF output.<br>d. Select reports from the left column and click **Add** to move each report to the report group on the right.<br><br> 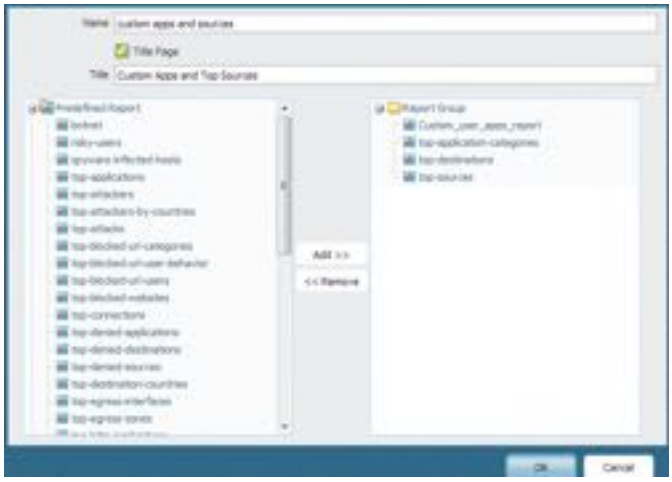 <br><br>The **Log View** report is a report type that is automatically created each time you create a custom report and uses the same name as the custom report. This report will show the logs that were used to build the contents of the custom report.<br><br>To include the log view data, when creating a report group, add your custom report under the **Custom Reports** list and then add the log view report by selecting the matching report name from the **Log View** list. The report will include the custom report data and the log data that was used to create the custom report.<br>e. Click **OK** to save the settings.<br>f. To use the report group, see Schedule Reports for Email Delivery. |

## Schedule Reports for Email Delivery

Reports can be scheduled for daily delivery or delivered weekly on a specified day. Scheduled reports are executed starting at 2:00 AM, and email delivery starts after all scheduled reports have been generated.

| Schedule Reports for Email Delivery | |
|---|---|
| Step 1 | Select **Monitor > PDF Reports > Email Scheduler** and click **Add**. |
| Step 2 | Enter a **Name** to identify the schedule. |
| Step 3 | Select the **Report Group** for email delivery. To set up a report group; see Manage Report Groups. |
| Step 4 | For the **Email Profile**, select an Email Server profile to use for delivering the reports, or click the **Email Profile** link to Create an Email server profile. |
| Step 5 | Select the frequency at which to generate and send the report in **Recurrence**. |
| Step 6 | The **Override Recipient email(s)** allows you to send this report exclusively to the recipients specified in this field. When you add recipients to the field, the report is not sent to the recipients configured in the email server profile. Use this option for those occasions when the report is for the attention of someone other than the administrators or recipients defined in the email server profile.<br> 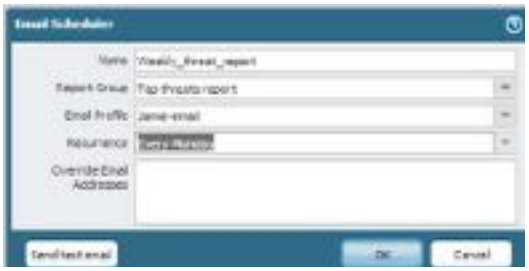 |

 © Palo Alto Networks, Inc.

# Use External Services for Monitoring

Using an external service to monitor the firewall enables you to receive alerts for important events, archive monitored information on systems with dedicated long-term storage, and integrate with third-party security monitoring tools. The following are some common scenarios for using external services:

- For immediate notification about important system events or threats, you can Monitor Device Statistics Using SNMP, Forward Traps to an SNMP Manager, or Configure Email Alerts.
- For long-term log storage and centralized firewall monitoring, you can Configure Syslog Monitoring to send log data to a syslog server. This enables integration with third-party security monitoring tools such as Splunk! or ArcSight.
- For monitoring statistics on the IP traffic that traverses firewall interfaces, you can Configure NetFlow Exports to view the statistics in a NetFlow collector.

You can Configure Log Forwarding from the firewalls directly to external services or from the firewalls to Panorama and then configure Panorama to forward logs to the servers. Refer to Log Forwarding Options for the factors to consider when deciding where to forward logs.

You can't aggregate NetFlow records on Panorama; you must send them directly from the firewalls to a NetFlow collector.

# Configure Log Forwarding

To use Panorama or Use External Services for Monitoring the firewall, you must configure the firewall to forward its logs. Before forwarding to external services, the firewall automatically converts the logs to the necessary format: syslog messages, SNMP traps, or email notifications. Before starting this procedure, ensure that Panorama or the external server that will receive the log data is already set up.

The PA-7000 Series firewall can't forward logs to Panorama, only to external services. However, when you use Panorama to monitor logs or generate reports for a device group that includes a PA-7000 Series firewall, Panorama queries the PA-7000 Series firewall in real-time to display its log data.

You can forward logs from the firewalls directly to external services or from the firewalls to Panorama and then configure Panorama to forward logs to the servers. Refer to Log Forwarding Options for the factors to consider when deciding where to forward logs.

You can use Secure Copy (SCP) commands from the CLI to export the entire log database to an SCP server and import it to another firewall. Because the log database is too large for an export or import to be practical on the PA-7000 Series firewall, it does not support these options. You can also use the web interface on all platforms to Schedule Log Exports to an SCP or FTP Server, but only on a per log type basis, not the entire log database.

| Configure Log Forwarding | | |
|---|---|---|
| Step 1 | Configure a server profile for each external service that will receive log data.<br><br>You can use separate profiles to send each log type to a different server. To increase availability, define multiple servers in a single profile. | • Create an Email server profile.<br>• Configure an SNMP Trap server profile. To enable the SNMP manager (trap server) to interpret firewall traps, you must load the Palo Alto Networks Supported MIBs into the SNMP manager and, if necessary, compile them. For details, refer to your SNMP management software documentation.<br>• Configure a Syslog server profile. If the syslog server requires client authentication, you must also Create a certificate to secure syslog communication over SSL. |
| Step 2 | Create a log forwarding profile.<br><br>The profile defines the destinations for Traffic, Threat, and WildFire Submission logs. (Threat logs include URL Filtering and Data Filtering logs.) | 1. Select **Objects > Log Forwarding** and click **Add**.<br>2. Enter a **Name** to identify the profile. If you want the firewall to automatically assign the profile to new security rules and zones, enter `default`. If you don't want a default profile, or you want to override an existing default profile, enter a **Name** that will help you identify the profile when assigning it to security rules and zones.<br><br>If no log forwarding profile named `default` exists, the profile selection is set to **None** by default in new security rules (**Log Forwarding** field) and new security zones (**Log Setting** field), although you can change the selection.<br>3. Perform the following steps for each log type and each severity level or WildFire verdict:<br>a. Select the **Panorama** check box if you want to aggregate firewall logs on Panorama. (You can then configure Panorama to forward the logs to external services.)<br>b. Select the **SNMP Trap**, **Email**, or **Syslog** server profile you configured for this log type, and click **OK**. |


© Palo Alto Networks, Inc.

| Configure Log Forwarding (Continued) | | |
|---|---|---|
| Step 3 | Assign the log forwarding profile to security rules.<br><br>To trigger log generation and forwarding, the rules require certain Security Profiles according to log type:<br>• Traffic logs—No security profile is necessary; the traffic only needs to match a specific security rule.<br>• Threat logs—The traffic must match any security profile assigned to a security rule.<br>• WildFire logs—The traffic must match a WildFire Analysis profile assigned to a security rule. | Perform the following steps for each rule that will trigger log forwarding:<br>1. Select **Policies > Security** and click the rule.<br>2. Select the **Actions** tab and select the **Log Forwarding** profile you just created.<br>3. In the **Profile Type** drop-down, select **Profiles** or **Group**, and then select the security profiles or **Group Profile** required to trigger log generation and forwarding.<br>4. For Traffic logs, select one or both of the **Log At Session Start** and **Log At Session End** check boxes, and click **OK**. |
| Step 4 | Configure the destinations of System, Config, HIP Match, and Correlation logs. | 1. Select **Device > Log Settings**.<br>2. Perform the following steps for each log type. For System and Correlation logs, start by clicking the Severity level. For Config and HIP Match logs, start by clicking the Edit icon.<br>a. Select the **Panorama** check box if you want to aggregate firewall logs on Panorama. You can then configure Panorama to forward the logs to the external services.<br>b. Select the **SNMP Trap**, **Email**, or **Syslog** server profile you configured for this log type and click **OK**. |
| Step 5 | (PA-7000 Series firewalls only) Configure a log card interface to perform log forwarding. | 1. Select **Network > Interfaces > Ethernet** and click **Add Interface**.<br>2. Select the **Slot** and **Interface Name**.<br>3. For the **Interface Type**, select **Log Card**.<br>4. Enter the **IP Address**, **Default Gateway**, and (for IPv4 only) **Netmask**.<br>5. Select **Advanced** and specify the **Link Speed**, **Link Duplex**, and **Link State**.<br>These fields default to **auto**, which specifies that the firewall automatically determines the values based on the connection. However, the minimum recommended **Link Speed** for any connection is **1000** (Mbps).<br>6. Click **OK** to save your changes. |

| Configure Log Forwarding (Continued) | |
|---|---|
| Step 6 Commit and verify your changes. | 1. Click **Commit** to complete the log forwarding configuration.<br>2. Verify the log destinations you configured are receiving firewall logs:<br>• Panorama—If the firewall forwards logs to an M-Series appliance, you must configure a Collector Group before Panorama will receive the logs. You can then verify log forwarding.<br>• Email server—Verify that the specified recipients are receiving logs as email notifications.<br>• Syslog server—Refer to the documentation for your syslog server to verify it is receiving logs as syslog messages.<br>• SNMP manager—Use an SNMP Manager to Explore MIBs and Objects to verify it is receiving logs as SNMP traps. |

 © Palo Alto Networks, Inc.

# Configure Email Alerts

You can configure email alerts for System, Config, HIP Match, Correlation, Threat, WildFire Submission, and Traffic logs.

| Configure Email Alerts | | |
|---|---|---|
| Step 1 | Create an Email server profile.<br> <br>You can use separate profiles to send email notifications for each log type to a different server. To increase availability, define multiple servers (up to four) in a single profile. | 1. Select **Device > Server Profiles > Email**.<br>2. Click **Add** and then enter a **Name** for the profile.<br>3. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available.<br>4. For each Simple Mail Transport Protocol (SMTP) server (email server), click **Add** and define the following information:<br>• **Name**—Name to identify the SMTP server (1-31 characters). This field is just a label and doesn't have to be the hostname of an existing email server.<br>• **Email Display Name**—The name to show in the From field of the email.<br>• **From**—The email address from which the Palo Alto Networks device sends emails.<br>• **To**—The email address to which the Palo Alto Networks device sends emails.<br>• **Additional Recipient**—If you want to send emails to a second account, enter the address here. You can add only one additional recipient. For multiple recipients, add the email address of a distribution list.<br>• **Email Gateway**—The IP address or hostname of the SMTP gateway to use for sending emails.<br>5. (Optional) Select the **Custom Log Format** tab and customize the format of the email messages. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide.<br>6. Click **OK** to save the Email server profile. |
| Step 2 | Configure email alerts for Traffic, Threat, and WildFire Submission logs. | 1. Create a log forwarding profile.<br>a. Select **Objects > Log Forwarding**, click **Add**, and enter a **Name** to identify the profile.<br>b. For each log type and each severity level or WildFire verdict, select the Email server profile and click **OK**.<br>2. Assign the log forwarding profile to security rules. |
| Step 3 | Configure email alerts for System, Config, HIP Match, and Correlation logs. | 1. Select **Device > Log Settings**.<br>2. For System and Correlation logs, click each Severity level, select the **Email** server profile, and click **OK**.<br>3. For Config and HIP Match logs, click the Edit icon, select the **Email** server profile, and click **OK**.<br>4. Click **Commit**. |

# Use Syslog for Monitoring

Syslog is a standard log transport mechanism that enables the aggregation of log data from different network devices—such as routers, firewalls, printers—from different vendors into a central repository for archiving, analysis, and reporting. Palo Alto Networks devices can forward every type of log they generate to an external syslog server. You can use TCP or SSL for reliable and secure log forwarding, or UDP for non-secure forwarding.


- Configure Syslog Monitoring
- Syslog Field Descriptions


 © Palo Alto Networks, Inc.

## Configure Syslog Monitoring

To Use Syslog for Monitoring a Palo Alto Networks device, create a Syslog server profile and assign it to the device log settings for each log type. Optionally, you can configure the header format used in syslog messages and enable client authentication for syslog over SSL.

| Configure Syslog Monitoring | | |
|---|---|---|
| Step 1 | Configure a Syslog server profile.<br> You can use separate profiles to send syslogs for each log type to a different server. To increase availability, define multiple servers (up to four) in a single profile. | 1. Select **Device > Server Profiles > Syslog**.<br>2. Click **Add** and enter a **Name** for the profile.<br>3. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available.<br>4. For each syslog server, click **Add** and enter the information that the firewall requires to connect to it:<br>• **Name**—Unique name for the server profile.<br>• **Server**—IP address or fully qualified domain name (FQDN) of the syslog server.<br>• **Transport**—Select TCP, UDP, or SSL as the method of communication with the syslog server.<br>• **Port**—The port number on which to send syslog messages (default is UDP on port 514); you must use the same port number on the firewall and the syslog server.<br>• **Format**—Select the syslog message format to use: **BSD** (the default) or **IETF**. Traditionally, **BSD** format is over UDP and **IETF** format is over TCP or SSL.<br>• **Facility**—Select a syslog standard value (default is **LOG_USER**) to calculate the priority (PRI) field in your syslog server implementation. Select the value that maps to how you use the PRI field to manage your syslog messages.<br>5. (Optional) To customize the format of the syslog messages that the firewall sends, select the **Custom Log Format** tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide.<br>6. Click **OK** to save the server profile. |
| Step 2 | Configure syslog forwarding for Traffic, Threat, and WildFire Submission logs. | 1. Create a log forwarding profile.<br>a. Select **Objects > Log Forwarding**, click **Add**, and enter a **Name** to identify the profile.<br>b. For each log type and each severity level or WildFire verdict, select the **Syslog** server profile and click **OK**.<br>2. Assign the log forwarding profile to security rules. |

| Configure Syslog Monitoring (Continued) | | |
|---|---|---|
| Step 3 | Configure syslog forwarding for System, Config, HIP Match, and Correlation logs. | 1. Select **Device > Log Settings**.<br>2. For System and Correlation logs, click each Severity level, select the **Syslog** server profile, and click **OK**.<br>3. For Config, HIP Match, and Correlation logs, click the Edit icon, select the **Syslog** server profile, and click **OK**. |
| Step 4 | (Optional) Configure the header format of syslog messages.<br><br>The log data includes the unique identifier of the device that generated the log. Choosing the header format provides more flexibility in filtering and reporting on the log data for some Security Information and Event Management (SIEM) servers.<br><br>This is a global setting and applies to all syslog server profiles configured on the device. | 1. Select **Device > Setup > Management** and edit the Logging and Reporting Settings.<br>2. Select the **Log Export and Reporting** tab and select the **Syslog HOSTNAME Format**:<br>• **FQDN** (default)—Concatenates the hostname and domain name defined on the sending device.<br>• **hostname**—Uses the hostname defined on the sending device.<br>• **ipv4-address**—Uses the IPv4 address of the device interface used to send logs. By default, this is the MGT interface.<br>• **ipv6-address**—Uses the IPv6 address of the device interface used to send logs. By default, this is the MGT interface.<br>• **none**—Leaves the hostname field unconfigured on the device. There is no identifier for the device that sent the logs.<br>3. Click **OK** to save your changes. |
| Step 5 | Create a certificate to secure syslog communication over SSL.<br><br>Required only if the syslog server uses client authentication. The syslog server uses the certificate to verify that the device is authorized to communicate with the syslog server.<br><br>Ensure the following conditions are met:<br>• The private key must be available on the sending device; the keys can't reside on a Hardware Security Module (HSM).<br>• The subject and the issuer for the certificate must not be identical.<br>• The syslog server and the sending device must have certificates that the same trusted certificate authority (CA) signed. Alternatively, you can generate a self-signed certificate on the device, export the certificate from the device, and import it in to the syslog server. | 1. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Generate**.<br>2. Enter a **Name** for the certificate.<br>3. In the **Common Name** field, enter the IP address of the device sending logs to the syslog server.<br>4. In **Signed by**, select the trusted CA or the self-signed CA that the syslog server and the sending device both trust.<br>The certificate can't be a **Certificate Authority** nor an **External Authority** (certificate signing request [CSR]).<br>5. Click **Generate**. The device generates the certificate and key pair.<br>6. Click the certificate Name to edit it, select the **Certificate for Secure Syslog** check box, and click **OK**. |

 © Palo Alto Networks, Inc.

| Configure Syslog Monitoring (Continued) | | |
|---|---|---|
| Step 6 | Commit your changes and review the logs on the syslog server. | 1. Click **Commit**.<br>2. To review the logs, refer to the documentation of your syslog management software. You can also review the Syslog Field Descriptions. |

# Syslog Field Descriptions

The following topics list the standard fields of each log type that Palo Alto Networks devices can forward to an external server, as well as the severity levels, custom formats, and escape sequences. To facilitate parsing, the delimiter is a comma: each field is a comma-separated value (CSV) string. The FUTURE_USE tag applies to fields that the devices do not currently implement.

WildFire Submission logs are a subtype of Threat log and use the same syslog format.


- Traffic Logs
- Threat Logs
- HIP Match Logs
- Config Logs
- System Logs
- Correlated Events (Logs)
- Syslog Severity
- Custom Log/Event Format
- Escape Sequences


## Traffic Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Ingress Interface, Egress Interface, Log Forwarding Profile, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Bytes, Bytes Sent, Bytes Received, Packets, Start Time, Elapsed Time, Category, FUTURE_USE, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Packets Sent, Packets Received, Session End Reason, Device Group Hierarchy Level 1*, Device Group Hierarchy Level 2*, Device Group Hierarchy Level 3*, Device Group Hierarchy Level 4*, Virtual System Name*, Device Name*, Action Source*

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Specifies type of log; values are traffic, threat, config, system and hip-match |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Subtype (subtype) | Subtype of traffic log; values are start, end, drop, and deny<br>• Start—session started<br>• End—session ended<br>• Drop—session dropped before the application is identified and there is no rule that allows the session.<br>• Deny—session dropped after the application is identified and there is a rule to block or no rule that allows the session. |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Source IP (src) | Original session source IP address |
| Destination IP (dst) | Original session destination IP address |
| NAT Source IP (natsrc) | If Source NAT performed, the post-NAT Source IP address |
| NAT Destination IP (natdst) | If Destination NAT performed, the post-NAT Destination IP address |
| Rule Name (rule) | Name of the rule that the session matched |
| Source User (srcuser) | Username of the user who initiated the session |
| Destination User (dstuser) | Username of the user to which the session was destined |
| Application (app) | Application associated with the session |
| Virtual System (vsys) | Virtual System associated with the session |
| Source Zone (from) | Zone the session was sourced from |
| Destination Zone (to) | Zone the session was destined to |
| Ingress Interface (inbound_if) | Interface that the session was sourced form |
| Egress Interface (outbound_if) | Interface that the session was destined to |
| Log Forwarding Profile (logset) | Log Forwarding Profile that was applied to the session |
| Session ID (sessionid) | An internal numerical identifier applied to each session |
| Repeat Count (repeatcnt) | Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only |
| Source Port (sport) | Source port utilized by the session |
| Destination Port (dport) | Destination port utilized by the session |
| NAT Source Port (natsport) | Post-NAT source port |
| NAT Destination Port (natdport) | Post-NAT destination port |

© Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Flags (flags) | 32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value:<br>• 0x80000000—session has a packet capture (PCAP)<br>• 0x02000000—IPv6 session<br>• 0x01000000—SSL session was decrypted (SSL Proxy)<br>• 0x00800000—session was denied via URL filtering<br>• 0x00400000—session has a NAT translation performed (NAT)<br>• 0x00200000—user information for the session was captured via the captive portal (Captive Portal)<br>• 0x00080000—X-Forwarded-For value from a proxy is in the source user field<br>• 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction)<br>• 0x00008000—session is a container page access (Container Page)<br>• 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above.<br>• 0x00000800—symmetric return was used to forward traffic for this session |
| Protocol (proto) | IP protocol associated with the session |
| Action (action) | Action taken for the session; possible values are:<br>• Allow—session was allowed by policy<br>• Deny—session was denied by policy<br>• Drop—session was dropped silently<br>• Drop ICMP—session was silently dropped with an ICMP unreachable message to the host or application<br>• Reset both—session was terminated and a TCP reset is sent to both the sides of the connection<br>• Reset client—session was terminated and a TCP reset is sent to the client<br>• Reset server—session was terminated and a TCP reset is sent to the server |
| Bytes (bytes) | Number of total bytes (transmit and receive) for the session |
| Bytes Sent (bytes_sent) | Number of bytes in the client-to-server direction of the session<br>Available on all models except the PA-4000 Series |
| Bytes Received (bytes_received) | Number of bytes in the server-to-client direction of the session<br>Available on all models except the PA-4000 Series |
| Packets (packets) | Number of total packets (transmit and receive) for the session |
| Start Time (start) | Time of session start |
| Elapsed Time (elapsed) | Elapsed time of the session |
| Category (category) | URL category associated with the session (if applicable) |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7000 Series firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama |
| Source Location (srcloc) | Source country or Internal region for private addresses; maximum length is 32 bytes |
| Destination Location (dstloc) | Destination country or Internal region for private addresses. Maximum length is 32 bytes |
| Packets Sent (pkts_sent) | Number of client-to-server packets for the session<br>Available on all models except the PA-4000 Series |
| Packets Received (pkts_received) | Number of server-to-client packets for the session<br>Available on all models except the PA-4000 Series |
| Session End Reason (session_end_reason) | The reason a session terminated. If the termination had multiple causes, this field displays only the highest priority reason. The possible session end reason values are as follows, in order of priority (where the first is highest):<br>• threat—The firewall detected a threat associated with a reset, drop, or block (IP address) action.<br>• policy-deny—The session matched a security rule with a deny or drop action.<br>• tcp-rst-from-client—The client sent a TCP reset to the server.<br>• tcp-rst-from-server—The server sent a TCP reset to the client.<br>• resources-unavailable—The session dropped because of a system resource limitation. For example, the session could have exceeded the number of out-of-order packets allowed per flow or the global out-of-order packet queue.<br>• tcp-fin—One host or both hosts in the connection sent a TCP FIN message to close the session.<br>• tcp-reuse—A session is reused and the firewall closes the previous session.<br>• decoder—The decoder detects a new connection within the protocol (such as HTTP-Proxy) and ends the previous connection.<br>• aged-out—The session aged out.<br>• unknown—This value applies in the following situations:<br>  • Session terminations that the preceding reasons do not cover (for example, a `clear session all` command).<br>  • For logs generated in a PAN-OS release that does not support the session end reason field (releases older than PAN-OS 6.1), the value will be `unknown` after an upgrade to the current PAN-OS release or after the logs are loaded onto the firewall.<br>  • In Panorama, logs received from firewalls for which the PAN-OS version does not support session end reasons will have a value of `unknown`.<br>• n/a—This value applies when the traffic log type is not `end`. |

| Field Name | Description |
|---|---|
| Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)<br>New in v7.0! | A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure.<br>If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods:<br>CLI command in configure mode: **`show readonly dg-meta-data`**<br>API query:<br>`/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>` |
| Virtual System Name (vsys_name)<br>New in v7.0! | The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems. |
| Device Name (device_name)<br>New in v7.0! | The hostname of the firewall on which the session was logged. |
| Action Source (action_source)<br>New in v7.0! | Specifies whether the action taken to allow or block an application was defined in the application or in policy. The actions can be allow, deny, drop, reset- server, reset-client or reset-both for the session. |

 © Palo Alto Networks, Inc.

## Threat Logs

**Format**: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Ingress Interface, Egress Interface, Log Forwarding Profile, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Miscellaneous, Threat ID, Category, Severity, Direction, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Content Type, PCAP_id, Filedigest, Cloud, URL Index*, User Agent, File Type, X-Forwarded-For, Referer, Sender, Subject, Recipient, Report ID, Device Group Hierarchy Level 1*, Device Group Hierarchy Level 2*, Device Group Hierarchy Level 3*, Device Group Hierarchy Level 4*, Virtual System Name*, Device Name*, FUTURE_USE,

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Specifies type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of threat log. Values include the following:<br>• data—Data pattern matching a Data Filtering profile.<br>• file—File type matching a File Blocking profile.<br>• flood—Flood detected via a Zone Protection profile.<br>• packet—Packet-based attack protection triggered by a Zone Protection profile.<br>• scan—Scan detected via a Zone Protection profile.<br>• spyware—Spyware detected via an Anti-Spyware profile.<br>• url—URL filtering log.<br>• virus—Virus detected via an Antivirus profile.<br>• vulnerability—Vulnerability exploit detected via a Vulnerability Protection profile.<br>• wildfire—A WildFire verdict generated when the firewall submits a file to WildFire per a WildFire Analysis profile and a verdict (malicious, grayware, or benign, depending on what you are logging) is logged in the WildFire Submissions log.<br>• wildfire-virus—Virus detected via an Antivirus profile. |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Source IP (src) | Original session source IP address |
| Destination IP (dst) | Original session destination IP address |
| NAT Source IP (natsrc) | If source NAT performed, the post-NAT source IP address |
| NAT Destination IP (natdst) | If destination NAT performed, the post-NAT destination IP address |
| Rule Name (rule) | Name of the rule that the session matched |
| Source User (srcuser) | Username of the user who initiated the session |

| Field Name | Description |
|---|---|
| Destination User (dstuser) | Username of the user to which the session was destined |
| Application (app) | Application associated with the session |
| Virtual System (vsys) | Virtual System associated with the session |
| Source Zone (from) | Zone the session was sourced from |
| Destination Zone (to) | Zone the session was destined to |
| Ingress Interface (inbound_if) | Interface that the session was sourced from |
| Egress Interface (outbound_if) | Interface that the session was destined to |
| Log Forwarding Profile (logset) | Log Forwarding Profile that was applied to the session |
| Session ID (sessionid) | An internal numerical identifier applied to each session |
| Repeat Count (repeatcnt) | Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only |
| Source Port (sport) | Source port utilized by the session |
| Destination Port (dport) | Destination port utilized by the session |
| NAT Source Port (natsport) | Post-NAT source port |
| NAT Destination Port (natdport) | Post-NAT destination port |
| Flags (flags) | 32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value:<br>• 0x80000000—session has a packet capture (PCAP)<br>• 0x02000000—IPv6 session<br>• 0x01000000—SSL session was decrypted (SSL Proxy)<br>• 0x00800000—session was denied via URL filtering<br>• 0x00400000—session has a NAT translation performed (NAT)<br>• 0x00200000—user information for the session was captured via the captive portal (Captive Portal)<br>• 0x00080000—X-Forwarded-For value from a proxy is in the source user field<br>• 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction)<br>• 0x00008000—session is a container page access (Container Page)<br>• 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above<br>• 0x00000800—symmetric return was used to forward traffic for this session |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Protocol (proto) | IP protocol associated with the session |
| Action (action) | Action taken for the session; values are alert, allow, deny, drop, drop-all-packets, reset-client, reset-server, reset-both, block-url.<br>• Alert—threat or URL detected but not blocked<br>• Allow—flood detection alert<br>• Deny—flood detection mechanism activated and deny traffic based on configuration<br>• Drop—threat detected and associated session was dropped<br>• Drop-all-packets—threat detected and session remains, but drops all packets<br>• Reset-client—threat detected and a TCP RST is sent to the client<br>• Reset-server—threat detected and a TCP RST is sent to the server<br>• Reset-both—threat detected and a TCP RST is sent to both the client and the server<br>• Block-url—URL request was blocked because it matched a URL category that was set to be blocked |
| Miscellaneous (misc) | Field with variable length with a maximum of 1023 characters<br>The actual URI when the subtype is URL<br>File name or file type when the subtype is file<br>File name when the subtype is virus<br>File name when the subtype is WildFire |
| Threat ID (threatid) | Palo Alto Networks identifier for the threat. It is a description string followed by a 64-bit numerical identifier in parentheses for some Subtypes:<br>• 8000 – 8099—scan detection<br>• 8500 – 8599—flood detection<br>• 9999—URL filtering log<br>• 10000 – 19999—sypware phone home detection<br>• 20000 – 29999—spyware download detection<br>• 30000 – 44999—vulnerability exploit detection<br>• 52000 – 52999—filetype detection<br>• 60000 – 69999—data filtering detection<br>• 100000 – 2999999—virus detection<br>• 3000000 – 3999999—WildFire signature feed<br>• 4000000-4999999—DNS Botnet signatures |
| Category (category) | For URL Subtype, it is the URL Category; For WildFire subtype, it is the verdict on the file and is either ‘malicious’, ‘grayware’, or ‘benign’; For other subtypes, the value is ‘any’. |
| Severity (severity) | Severity associated with the threat; values are informational, low, medium, high, critical |

| Field Name | Description |
|---|---|
| Direction (direction) | Indicates the direction of the attack, client-to-server or server-to-client:<br>• 0—direction of the threat is client to server<br>• 1—direction of the threat is server to client |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space. This field is not supported on PA-7000 Series firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama. |
| Source Location (srcloc) | Source country or Internal region for private addresses. Maximum length is 32 bytes. |
| Destination Location (dstloc) | Destination country or Internal region for private addresses. Maximum length is 32 bytes. |
| Content Type (contenttype) | Applicable only when Subtype is URL.<br>Content type of the HTTP response data. Maximum length 32 bytes. |
| PCAP ID (pcap_id) | The packet capture (pcap) ID is a 64 bit unsigned integral denoting an ID to correlate threat pcap files with extended pcaps taken as a part of that flow. All threat logs will contain either a pcap_id of 0 (no associated pcap), or an ID referencing the extended pcap file. |
| File Digest (filedigest) | Only for WildFire subtype; all other types do not use this field<br>The filedigest string shows the binary hash of the file sent to be analyzed by the WildFire service. |
| Cloud (cloud) | Only for WildFire subtype; all other types do not use this field.<br>The cloud string displays the FQDN of either the WildFire appliance (private) or the WildFire cloud (public) from where the file was uploaded for analysis. |
| URL Index (url_idx)<br>New in v7.0! | Used in URL Filtering and WildFire subtypes.<br>When an application uses TCP keepalives to keep a connection open for a length of time, all the log entries for that session have a single session ID. In such cases, when you have a single threat log (and session ID) that includes multiple URL entries, the url_idx is a counter that allows you to correlate the order of each log entry within the single session.<br>For example, to learn the URL of a file that the firewall forwarded to WildFire for analysis, locate the session ID and the url_idx from the WildFire Submissions log and search for the same session ID and url_idx in your URL filtering logs. The log entry that matches the session ID and url_idx will contain the URL of the file that was forwarded to WildFire. |
| User Agent (user_agent) | Only for the URL Filtering subtype; all other types do not use this field.<br>The User Agent field specifies the web browser that the user used to access the URL, for example Internet Explorer. This information is sent in the HTTP request to the server. |
| File Type (filetype) | Only for WildFire subtype; all other types do not use this field.<br>Specifies the type of file that the firewall forwarded for WildFire analysis. |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| X-Forwarded-For (xff) | Only for the URL Filtering subtype; all other types do not use this field.<br>The X-Forwarded-For field in the HTTP header contains the IP address of the user who requested the web page. It allows you to identify the IP address of the user, which is useful particularly if you have a proxy server on your network that replaces the user IP address with its own address in the source IP address field of the packet header. |
| Referer (referer) | Only for the URL Filtering subtype; all other types do not use this field.<br>The Referer field in the HTTP header contains the URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested. |
| Sender (sender) | Only for WildFire subtype; all other types do not use this field.<br>Specifies the name of the sender of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall. |
| Subject (subject) | Only for WildFire subtype; all other types do not use this field.<br>Specifies the subject of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall. |
| Recipient (recipient) | Only for WildFire subtype; all other types do not use this field.<br>Specifies the name of the receiver of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall. |
| Report ID (reportid) | Only for WildFire subtype; all other types do not use this field.<br>Identifies the analysis request on the WildFire cloud or the WildFire appliance. |
| Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)<br>New in v7.0! | A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure.<br>If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods:<br>CLI command in configure mode: **`show readonly dg-meta-data`**<br>API query:<br>`/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>` |
| Virtual System Name (vsys_name)<br>New in v7.0! | The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems. |
| Device Name (device_name)<br>New in v7.0! | The hostname of the firewall on which the session was logged. |

## HIP Match Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source User, Virtual System, Machine name, OS, Source Address, HIP, Repeat Count, HIP Type, FUTURE_USE, FUTURE_USE, Sequence Number, Action Flags, Device Group Hierarchy Level 1*, Device Group Hierarchy Level 2*, Device Group Hierarchy Level 3*, Device Group Hierarchy Level 4*, Virtual System Name*, Device Name*

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of HIP match log; unused |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Source User (srcuser) | Username of the user who initiated the session |
| Virtual System (vsys) | Virtual System associated with the HIP match log |
| Machine Name (machinename) | Name of the user's machine |
| OS | The operating system installed on the user's machine or device (or on the client system) |
| Source Address (src) | IP address of the source user |
| HIP (matchname) | Name of the HIP object or profile |
| Repeat Count (repeatcnt) | Number of times the HIP profile matched |
| HIP Type (matchtype) | Whether the hip field represents a HIP object or a HIP profile |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7000 Series firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama |
| Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4) New in v7.0! | A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure.<br>If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods:<br>CLI command in configure mode: **`show readonly dg-meta-data`**<br>API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>` |

 © Palo Alto Networks, Inc.

| Field Name | Description |
|---|---|
| Virtual System Name (vsys_name)<br>New in v7.0! | The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems. |
| Device Name (device_name)<br>New in v7.0! | The hostname of the firewall on which the session was logged. |

## Config Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Host, Virtual System, Command, Admin, Client, Result, Configuration Path, Sequence Number, Action Flags, Before Change Detail, After Change Detail, Device Group Hierarchy Level 1*, Device Group Hierarchy Level 2*, Device Group Hierarchy Level 3*, Device Group Hierarchy Level 4*, Virtual System Name*, Device Name*

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of configuration log; unused |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Host (host) | Hostname or IP address of the client machine |
| Virtual System (vsys) | Virtual System associated with the configuration log |
| Command (cmd) | Command performed by the Admin; values are add, clone, commit, delete, edit, move, rename, set. |
| Admin (admin) | Username of the Administrator performing the configuration |
| Client (client) | Client used by the Administrator; values are Web and CLI |
| Result (result) | Result of the configuration action; values are Submitted, Succeeded, Failed, and Unauthorized |
| Configuration Path (path) | The path of the configuration command issued; up to 512 bytes in length |
| Sequence Number (seqno) | A 64bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7000 Series firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama. |
| Before Change Detail (before_change_detail) | This field is in custom logs only; it is not in the default format.<br>It contains the full xpath before the configuration change. |

| Field Name | Description |
|---|---|
| After Change Detail (after_change_detail) | This field is in custom logs only; it is not in the default format.<br>It contains the full xpath after the configuration change. |
| Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)<br>New in v7.0! | A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure.<br>If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods:<br>CLI command in configure mode: **`show readonly dg-meta-data`**<br>API query:<br>`/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>` |
| Virtual System Name (vsys_name)<br>New in v7.0! | The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems. |
| Device Name (device_name)<br>New in v7.0! | The hostname of the firewall on which the session was logged. |

 © Palo Alto Networks, Inc.

## System Logs

**Format:** FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Virtual System, Event ID, Object, FUTURE_USE, FUTURE_USE, Module, Severity, Description, Sequence Number, Action Flags, Device Group Hierarchy Level 1*, Device Group Hierarchy Level 2*, Device Group Hierarchy Level 3*, Device Group Hierarchy Level 4*, Virtual System Name*, Device Name*

| Field Name | Description |
|---|---|
| Receive Time (receive_time) | Time the log was received at the management plane |
| Serial Number (serial) | Serial number of the device that generated the log |
| Type (type) | Type of log; values are traffic, threat, config, system and hip-match |
| Subtype (subtype) | Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn |
| Generated Time (time_generated) | Time the log was generated on the dataplane |
| Virtual System (vsys) | Virtual System associated with the configuration log |
| Event ID (eventid) | String showing the name of the event |
| Object (object) | Name of the object associated with the system event |
| Module (module) | This field is valid only when the value of the Subtype field is general. It provides additional information about the sub-system generating the log; values are general, management, auth, ha, upgrade, chassis |
| Severity (severity) | Severity associated with the event; values are informational, low, medium, high, critical |
| Description (opaque) | Detailed description of the event, up to a maximum of 512 bytes |
| Sequence Number (seqno) | A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7000 Series firewalls. |
| Action Flags (actionflags) | A bit field indicating if the log was forwarded to Panorama |
| Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4) New in v7.0! | A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure.<br>If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods:<br>CLI command in configure mode: **`show readonly dg-meta-data`**<br>API query:<br>`/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>` |

| Field Name | Description |
|---|---|
| Virtual System Name (vsys_name)<br>New in v7.0! | The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems. |
| Device Name (device_name)<br>New in v7.0! | The hostname of the firewall on which the session was logged. |

 © Palo Alto Networks, Inc.

## Correlated Events (Logs)

**Format:** Receive Time*, Serial Number*, Device Name*, Type*, Virtual System Name*, Virtual System ID*, Device Group Hierarchy Level 1*, Device Group Hierarchy Level 2*, Device Group Hierarchy Level 3*, Device Group Hierarchy Level 4*, Source User*, Source*, Object Name*, Object ID*, Category*, Severity*, Evidence*

*All these fields are new in v7.0.5!

| Field Name | Description |
|---|---|
| Receive Time (receive_time)<br>*New in v7.0! | The time the firewall or Panorama recorded the event match. |
| Serial Number (serial)<br>*New in v7.0! | Serial number of the firewall that generated the log. |
| Device Name (device_name)<br>*New in v7.0! | The hostname of the firewall that generated the log. |
| Type (type)<br>*New in v7.0! | Correlation log |
| Virtual System (vsys)<br>*New in v7.0! | Name of the virtual system on the firewall that generated the correlation log.<br>(not pertinent to Panorama) |
| Virtual System ID(vsys_id)<br>*New in v7.0! | ID of the virtual system on the firewall that generated the correlation log.<br>(not pertinent to Panorama) |
| Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)<br>*New in v7.0! | A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure.<br>If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods:<br>CLI command in configure mode: **`show readonly dg-meta-data`**<br>API query:<br>`/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>` |
| Source User (srcuser)<br>*New in v7.0! | Username of the user who triggered the event. |
| Source (src)<br>*New in v7.0! | IP address of the user who triggered the event. |
| Object Name (object_name)<br>*New in v7.0! | Name of the correlation object that was matched on. For example, Beacon Detection. |

| Field Name | Description |
|---|---|
| Object ID (object_id)<br>*New in v7.0! | ID of the correlation object associated with the event. |
| Category (category)<br>*New in v7.0! | Category of the correlation object that triggered the match. For example, Compromised Host. |
| Severity (severity)<br>*New in v7.0! | Severity associated with the event; values are informational, low, medium, high, critical |
| Evidence (evidence)<br>*New in v7.0! | A summary statement with evidence to indicate how many times the host has matched against the conditions defined in the correlation object. For example, Host visited known malware URL (19 times). |

## Syslog Severity

The syslog severity is set based on the log type and contents.

| Log Type/Severity | Syslog Severity |
|---|---|
| Traffic | Info |
| Config | Info |
| Threat/System—Informational | Info |
| Threat/System—Low | Notice |
| Threat/System—Medium | Warning |
| Threat/System—High | Error |
| Threat/System—Critical | Critical |

## Custom Log/Event Format

To facilitate the integration with external log parsing systems, the firewall allows you to customize the log format; it also allows you to add custom *Key: Value* attribute pairs. Custom message formats can be configured under **Device > Server Profiles > Syslog > Syslog Server Profile > Custom Log Format**.

To achieve ArcSight Common Event Format (CEF) compliant log formatting, refer to the CEF Configuration Guide.

## Escape Sequences

Any field that contains a comma or a double-quote is enclosed in double quotes. Furthermore, if a double-quote appears inside a field it is escaped by preceding it with another double-quote. To maintain backward compatibility, the Misc field in threat log is always enclosed in double-quotes.

 © Palo Alto Networks, Inc.

# SNMP Monitoring and Traps

The following topics describe how Palo Alto Networks devices implement Simple Network Management Protocol (SNMP), and the procedures to configure SNMP monitoring and trap delivery.


- SNMP for Palo Alto Networks Devices
- Use an SNMP Manager to Explore MIBs and Objects
- Enable SNMP Services for Firewall-Secured Network Elements
- Monitor Device Statistics Using SNMP
- Forward Traps to an SNMP Manager
- Supported MIBs


© Palo Alto Networks, Inc.

## SNMP for Palo Alto Networks Devices

You can use a Simple Network Management Protocol (SNMP) manager to monitor event-driven alerts and operational statistics for Palo Alto Networks devices and the traffic that those devices secure. The statistics and traps can help you identify resource limitations, system changes or failures, and malware attacks. You configure devices to send the alerts by forwarding log data as traps, and enable devices to send statistics in response to GET messages (requests) from your SNMP manager. Each trap and statistic has an object identifier (OID). Related OIDs are organized hierarchically within the Management Information Bases (MIBs) that you load into the SNMP manager to enable monitoring.

Palo Alto Networks devices support SNMP Version 2c and Version 3. Decide which to use based on the version that other devices in your network support and on your network security requirements. SNMPv3 is more secure and enables more granular access control for device statistics than SNMPv2c. The following table summarizes the security features of each version. You select the version and configure the security features when you Monitor Device Statistics Using SNMP and Forward Traps to an SNMP Manager.

| SNMP Version | Device/User Authentication | Message Privacy | Message Integrity | MIB Access Granularity |
|---|---|---|---|---|
| SNMPv2c | Community string | No (cleartext) | No | SNMP community access for all MIBs on a device |
| SNMPv3 | EngineID, username, and authentication password (SHA hashing for the password) | Privacy password for AES 128 encryption of SNMP messages | Yes | User access based on views that include or exclude specific OIDs |

Figure: SNMP for Palo Alto Networks Devices illustrates a deployment in which firewalls forward traps to an SNMP manager while also forwarding logs to Log Collectors. Alternatively, you could configure the Log Collectors to forward the firewall traps to the SNMP manager. For details on these deployments, refer to Log Forwarding Options. In all deployments, the SNMP manager gets statistics directly from the devices. In this example, a single SNMP manager collects both traps and statistics, though you can use separate managers for these functions if that better suits your network.

 © Palo Alto Networks, Inc.

**Figure: SNMP for Palo Alto Networks Devices**

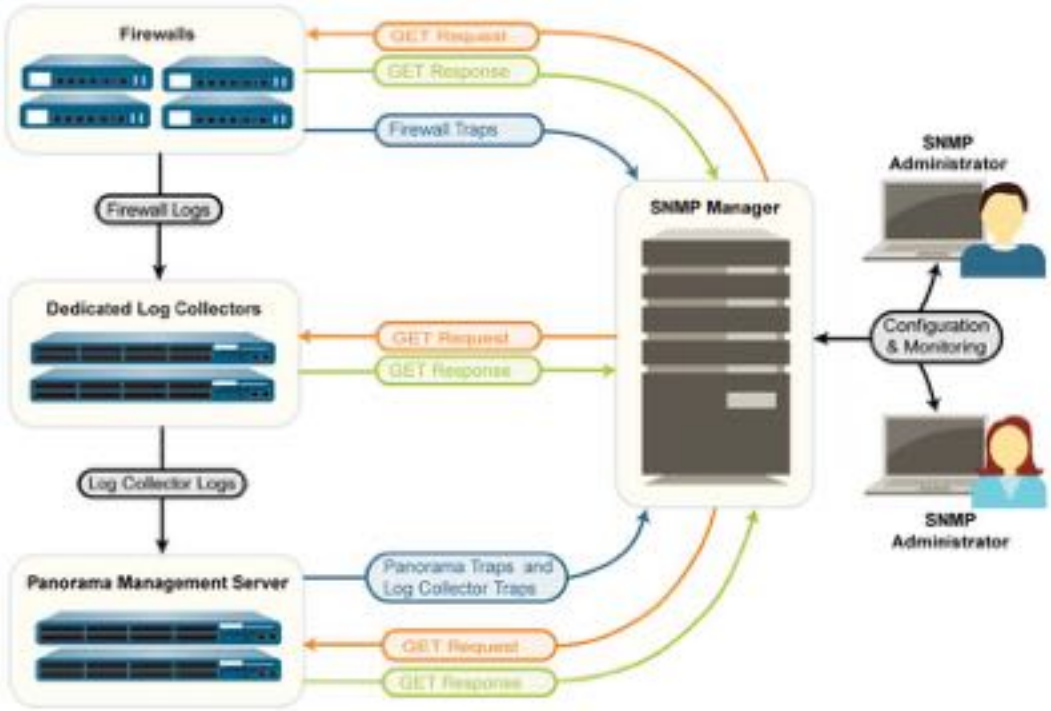

# Use an SNMP Manager to Explore MIBs and Objects

To use SNMP for monitoring Palo Alto Networks devices, you must first load the Supported MIBs into your SNMP manager and determine which object identifiers (OIDs) correspond to the statistics and traps you want to monitor. The following topics provide an overview of how to find OIDs and MIBs in an SNMP manager. For the specific steps to perform these tasks, refer to your SNMP management software.

- Identify a MIB Containing a Known OID
- Walk a MIB
- Identify the OID for a Palo Alto Networks Device Statistic or Trap

## Identify a MIB Containing a Known OID

If you already know the OID for a particular SNMP object (device statistic or trap) and want to know the OIDs of similar objects so you can monitor them, you can explore the MIB that contains the known OID.

| **Identify a MIB Containing a Known OID** | |
|---|---|
| Step 1 | Load all the Supported MIBs into your SNMP manager. |

 © Palo Alto Networks, Inc.

**Identify a MIB Containing a Known OID (Continued)**

Step 2 Search the entire MIB tree for the known OID. The search result displays the MIB path for the OID, as well as information about the OID (for example, name, status, and description). You can then select other OIDs in the same MIB to see information about them.

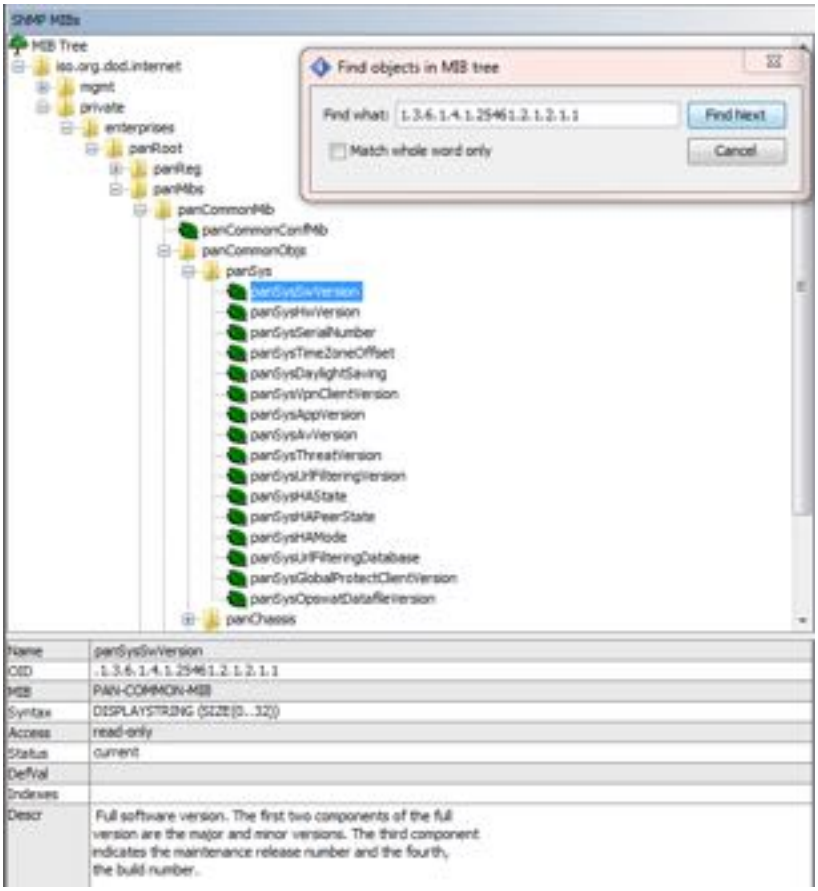


Step 3 Optionally, Walk a MIB to display all its objects.

## Walk a MIB

If you want to see which SNMP objects (device statistics and traps) are available for monitoring, displaying all the objects of a particular MIB can be useful. To do this, load the Supported MIBs into your SNMP manager and perform a *walk* on the desired MIB. To list the traps that Palo Alto Networks devices support, walk the panCommonEventEventsV2 MIB. In the following example, walking the PAN-COMMON-MIB.my displays the following list of OIDs and their values for certain device statistics:

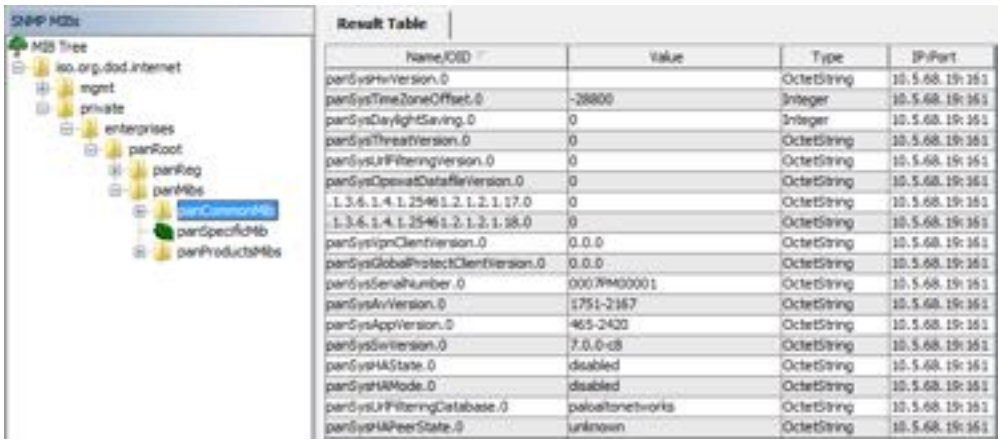


## Identify the OID for a Palo Alto Networks Device Statistic or Trap

To use an SNMP manager for monitoring Palo Alto Networks devices, you must know the OIDs of the device statistics and traps you want to monitor.

**Identify the OID for a Palo Alto Networks Device Statistic or Trap**

Step 1 Review the Supported MIBs to determine which one contains the type of statistic you want. For example, the PAN-COMMON-MIB.my contains device version information. The panCommonEventEventsV2 MIB contains all the traps that Palo Alto Networks devices support.

Step 2 Open the MIB in a text editor and perform a keyword search. For example, using **`Hardware version`** as a search string in PAN-COMMON-MIB identifies the panSysHwVersion object:

```
panSysHwVersion OBJECT-TYPE
     SYNTAX      DisplayString (SIZE(0..128))
     MAX-ACCESS  read-only
     STATUS      current
     DESCRIPTION
          "Hardware version of the unit."
     ::= {panSys 2}
```

 © Palo Alto Networks, Inc.

**Identify the OID for a Palo Alto Networks Device Statistic or Trap (Continued)**

Step 3 In a MIB browser, search the MIB tree for the identified object name to display its OID. For example, the panSysHwVersion object has an OID of 1.3.6.1.4.1.25461.2.1.2.1.2.

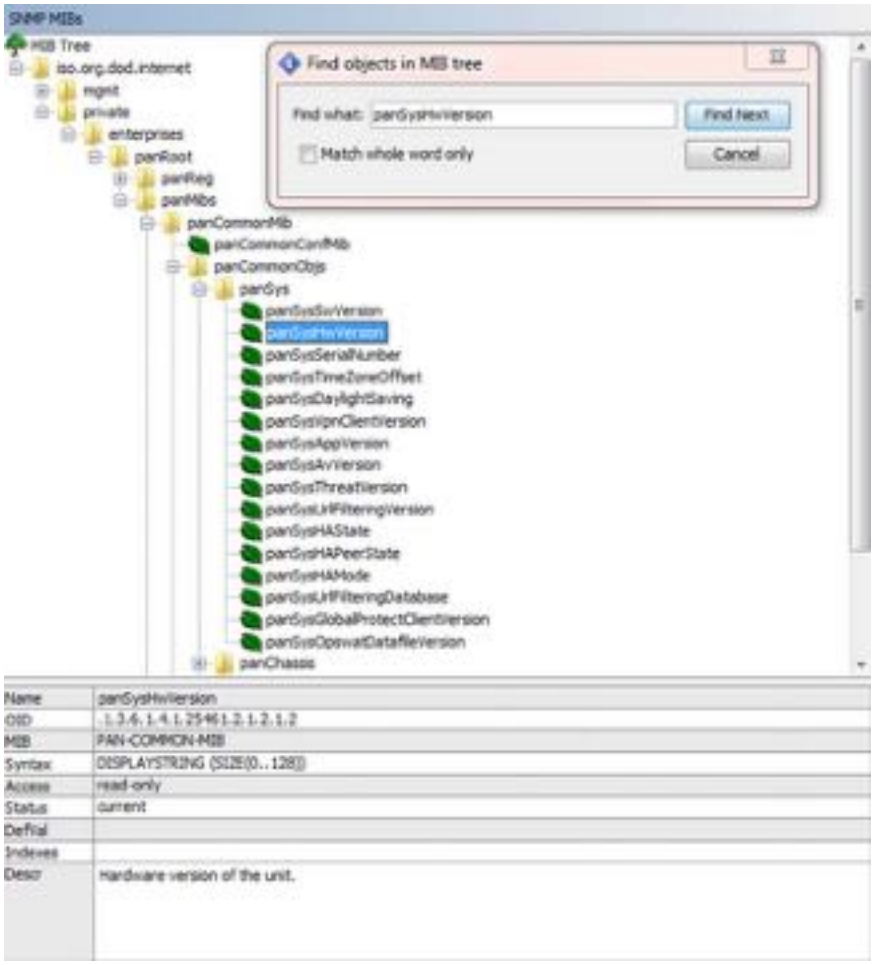

## Enable SNMP Services for Firewall-Secured Network Elements

If you will use Simple Network Management Protocol (SNMP) to monitor or manage network elements (for example, switches and routers) that are within the security zones of Palo Alto Networks firewalls, you must create a security rule that allows SNMP services for those elements.

You don't need a security rule to enable SNMP monitoring of Palo Alto Networks devices. For details, see Monitor Device Statistics Using SNMP.

| Enable SNMP Services for Firewall-Secured Network Elements | | |
|---|---|---|
| Step 1 | Create an application group. | 1. Select **Objects > Application Group** and click **Add**.<br>2. Enter a **Name** to identify the application group.<br>3. Click **Add**, type `snmp`, and select **snmp** and **snmp-trap** from the drop-down.<br>4. Click **OK** to save the application group. |
| Step 2 | Create a security rule to allow SNMP services. | 1. Select **Policies > Security** and click **Add**.<br>2. In the **General** tab, enter a **Name** for the rule.<br>3. In the **Source** and **Destination** tabs, click **Add** and enter a **Source Zone** and a **Destination Zone** for the traffic.<br>4. In the **Applications** tab, click **Add**, type the name of the applications group you just created, and select it from the drop-down.<br>5. In the **Actions** tab, verify that the **Action** is set to **Allow**, and then click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# Monitor Device Statistics Using SNMP

The statistics that a Simple Network Management Protocol (SNMP) manager collects from Palo Alto Networks devices can help you gauge the health of your network (devices and connections), identify resource limitations, and monitor traffic or processing loads. The statistics include information such as interface states (up or down), active user sessions, concurrent sessions, session utilization, temperature, and system uptime.

You can't configure an SNMP manager to control Palo Alto Networks devices (using SET messages), only to collect statistics from them (using GET messages).

For details on how SNMP is implemented for Palo Alto Networks devices, see SNMP for Palo Alto Networks Devices.

| Monitor Device Statistics Using SNMP | | |
|---|---|---|
| Step 1 | Configure the SNMP Manager to get statistics from devices. | The following steps provide an overview of the tasks you perform on the SNMP manager. For the specific steps, refer to the documentation of your SNMP manager.<br>1. To enable the SNMP manager to interpret device statistics, load the Supported MIBs for Palo Alto Networks devices and, if necessary, compile them.<br>2. For each device that the SNMP manager will monitor, define the connection settings (IP address and port) and authentication settings (SNMPv2c community string or SNMPv3 EngineID/username/password) for the device. Note that all Palo Alto Networks devices use port 161.<br>The SNMP manager can use the same or different connection and authentication settings for multiple devices. The settings must match those you define when you configure SNMP on the device (see Step 3). For example, if you use SNMPv2c, the community string you define when configuring the device must match the community string you define in the SNMP manager for that device.<br>3. Determine the object identifiers (OIDs) of the statistics you want to monitor. For example, to monitor the session utilization percentage of a firewall, a MIB browser shows that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.1.2.3.1.0 in PAN-COMMON-MIB.my. For details, see Use an SNMP Manager to Explore MIBs and Objects.<br>4. Configure the SNMP manager to monitor the desired OIDs. |

| Monitor Device Statistics Using SNMP (Continued) | |
|---|---|
| Step 2 Enable SNMP traffic on a device interface.<br><br>This is the interface that will receive statistics requests from the SNMP manager.<br><br>PAN-OS doesn't synchronize management (MGT) interface settings for devices in a high availability (HA) configuration. You must configure the interface for each HA peer. | Perform this step in the device web interface.<br><br>To enable SNMP traffic on the MGT interface:<br>1. Select **Device > Setup > Management** and edit the Management Interface Settings.<br>2. In the **Services** section, select the **SNMP** check box.<br>3. Click **OK** and **Commit**.<br><br>To enable SNMP traffic on any other interface:<br>1. Create an interface management profile for SNMP services:<br>a. Select **Network > Network Profiles > Interface Mgmt** and click **Add**.<br>b. Enter a **Name** for the profile, then select the check boxes for **SNMP** and any other services the interface must support.<br>c. Click **OK** to save the profile.<br>2. Assign the profile to the interface that will receive the SNMP requests:<br>a. Select **Network > Interfaces** and **Add** or edit the interface that will receive the SNMP requests. The interface type must be Layer 3 Ethernet.<br>b. Select **Advanced > Other Info** and select the **Management Profile** you just created.<br>c. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

| Monitor Device Statistics Using SNMP (Continued) | | |
|---|---|---|
| Step 3 | Configure the device to respond to statistics requests from an SNMP manager.<br><br>PAN-OS doesn't synchronize SNMP response settings for devices in a high availability (HA) configuration. You must configure these settings for each HA peer. | 1. Select **Device > Setup > Operations** and, in the Miscellaneous section, click **SNMP Setup**.<br>2. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP for Palo Alto Networks Devices.<br>• **V2c**—Enter the **SNMP Community String**, which identifies a community of SNMP managers and monitored devices, and serves as a password to authenticate the community members to each other.<br>As a best practice, don't use the default community string `public`; it's well known and therefore not secure.<br>• **V3**—Create at least one SNMP view group and one user. User accounts and views provide authentication, privacy, and access control when devices forward traps and SNMP managers get device statistics.<br>– **Views**—Each view is a paired OID and bitwise mask: the OID specifies a MIB and the mask (in hexadecimal format) specifies which objects are accessible within (include matching) or outside (exclude matching) that MIB. Click **Add** in the first list and enter a **Name** for the group of views. For each view in the group, click **Add** and configure the view **Name**, **OID**, matching **Option** (**include** or **exclude**), and **Mask**.<br>– **Users**: Click **Add** in the second list, enter a username under **Users**, select the **View** group from the drop-down, enter the authentication password (**Auth Password**) used to authenticate to the SNMP manager, and enter the privacy password (**Priv Password**) used to encrypt SNMP messages to the SNMP manager.<br>3. Click **OK** and **Commit**. |
| Step 4 | Monitor the firewall statistics in an SNMP manager. | Refer to the documentation of your SNMP manager.<br><br>When monitoring statistics related to firewall interfaces, you must match the interface indexes in the SNMP manager with interface names in the firewall web interface. For details, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors. |

# Forward Traps to an SNMP Manager

Simple Network Management Protocol (SNMP) traps can alert you to system events (failures or changes in hardware or software of Palo Alto Networks devices) or to threats (traffic that matches a firewall security rule) that require immediate attention.

To see the list of traps that Palo Alto Networks devices support, use your SNMP Manager to access the panCommonEventEventsV2 MIB. For details, see Use an SNMP Manager to Explore MIBs and Objects.

For details on how for Palo Alto Networks devices implement SNMP, see SNMP for Palo Alto Networks Devices.

| Forward Firewall Traps to an SNMP Manager | | |
|---|---|---|
| Step 1 | Enable the SNMP manager to interpret the traps it receives. | Load the Supported MIBs for Palo Alto Networks devices and, if necessary, compile them. For the specific steps, refer to the documentation of your SNMP manager. |
| Step 2 | Configure an SNMP Trap server profile.<br><br>The profile defines how the device accesses the SNMP managers (trap servers). You can define up to four SNMP managers for each profile.<br><br>Optionally, you can configure separate SNMP Trap server profiles for different log types, severity levels, and WildFire verdicts. | 1. Log in to the web interface of the Palo Alto Networks device.<br>2. Select **Device > Server Profiles > SNMP Trap**.<br>3. Click **Add** and enter a **Name** for the profile.<br>4. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available.<br>5. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP for Palo Alto Networks Devices.<br>• **V2c**—For each server, click **Add** and enter the server **Name**, IP address (**SNMP Manager**), and **Community String**. The community string identifies a community of SNMP managers and monitored devices, and serves as a password to authenticate the community members to each other.<br>As a best practice, don't use the default community string `public`; it's well known and therefore not secure.<br>• **V3**—For each server, click **Add** and enter the server **Name**, IP address (**SNMP Manager**), SNMP **User** account (this must match a username defined in the SNMP manager), **EngineID** used to uniquely identify the device (you can leave the field blank to use the device serial number), authentication password (**Auth Password**) used to authenticate to the server, and privacy password (**Priv Password**) used to encrypt SNMP messages to the server.<br>6. Click **OK** to save the server profile. |

 © Palo Alto Networks, Inc.

| Forward Firewall Traps to an SNMP Manager (Continued) | |
|---|---|
| Step 3 Configure log forwarding. | 1. Configure the destinations of Traffic, Threat, and WildFire traps:<br>a. Create a log forwarding profile. For each log type and each severity level or WildFire verdict, select the **SNMP Trap** server profile.<br>b. Assign the log forwarding profile to security rules. The rules will trigger trap generation and forwarding.<br>2. Configure the destinations of System, Config, HIP Match, and Correlation logs. For each log (trap) type and severity level, select the **SNMP Trap** server profile.<br>3. Click **Commit**. |
| Step 4 Monitor the traps in an SNMP manager. | Refer to the documentation of your SNMP manager.<br> When monitoring traps related to firewall interfaces, you must match the interface indexes in the SNMP manager with interface names in the firewall web interface. For details, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors. |

## Supported MIBs

The following table lists the Simple Network Management Protocol (SNMP) management information bases (MIBs) that Palo Alto Networks devices support. You must load these MIBs into your SNMP manager to monitor the objects (device statistics and traps) that are defined in the MIBs. For details, see Use an SNMP Manager to Explore MIBs and Objects.

| MIB Type | Supported MIBs |
|---|---|
| **Standard**—The Internet Engineering Task Force (IETF) maintains most standard MIBs. You can download the MIBs from the IETF website.<br>Palo Alto Networks devices don't support every object (OID) in every one of these MIBs. See the Supported MIBs links for an overview of the supported OIDs. | MIB-II<br>IF-MIB<br>HOST-RESOURCES-MIB<br>ENTITY-MIB<br>ENTITY-SENSOR-MIB<br>ENTITY-STATE-MIB<br>IEEE 802.3 LAG MIB<br>LLDP-V2-MIB.my |
| **Enterprise**—You can download the enterprise MIBs from the Palo Alto Networks Technical Documentation site. | PAN-COMMON-MIB.my<br>PAN-GLOBAL-REG-MIB.my<br>PAN-GLOBAL-TC-MIB.my<br>PAN-LC-MIB.my<br>PAN-PRODUCT-MIB.my<br>PAN-ENTITY-EXT-MIB.my<br>PAN-TRAPS.my |

### MIB-II

MIB-II provides object identifiers (OIDs) for network management protocols in TCP/IP-based networks. Use this MIB to monitor general information about devices and interfaces. For example, you can analyze trends in bandwidth usage by interface type (ifType object) to determine if the firewall needs more interfaces of that type to accommodate spikes in traffic volume.

Palo Alto Networks devices support only the following object groups:

| Object Group | Description |
|---|---|
| system | Provides device information such as the hardware model, system uptime, FQDN, and physical location. |
| interfaces | Provides statistics for physical and logical interfaces such as type, current bandwidth (speed), operational status (for example, up or down), and discarded packets. Logical interface support includes VPN tunnels, aggregate groups, Layer 2 subinterfaces, Layer 3 subinterfaces, loopback interfaces, and VLAN interfaces. |

RFC 1213 defines this MIB.

 © Palo Alto Networks, Inc.

## IF-MIB

IF-MIB supports interface types (physical and logical) and larger counters (64K) beyond those defined in MIB-II. Use this MIB to monitor interface statistics in addition to those that MIB-II provides. For example, to monitor the current bandwidth of high-speed interfaces (greater than 2.2Gps) such as the 10G interfaces of the PA-5000 Series firewalls, you must check the ifHighSpeed object in IF-MIB instead of the ifSpeed object in MIB-II. IF-MIB statistics can be useful when evaluating the capacity of your network.

Palo Alto Networks devices support only the ifXTable in IF-MIB, which provides interface information such as the number of multicast and broadcast packets transmitted and received, whether an interface is in promiscuous mode, and whether an interface has a physical connector.

RFC 2863 defines this MIB.

## HOST-RESOURCES-MIB

HOST-RESOURCES-MIB provides information for host computer resources. Use this MIB to monitor CPU and memory usage statistics for devices. For example, checking the current CPU load (hrProcessorLoad object) can help you troubleshoot performance issues on the firewall.

Palo Alto Networks devices support portions of the following object groups:

| Object Group | Description |
|---|---|
| hrDevice | Provides information such as CPU load, storage capacity, and partition size. The hrProcessorLoad OIDs provide an average of the cores that process packets. For the PA-5060 firewall, which has multiple dataplanes (DPs), the average is of the cores across all the three DPs that process packets. |
| hrSystem | Provides information such as device uptime, number of current user sessions, and number of current processes. |
| hrStorage | Provides information such as the amount of used storage. |

RFC 2790 defines this MIB.

## ENTITY-MIB

ENTITY-MIB provides OIDs for multiple logical and physical components. Use this MIB to determine what physical components are loaded on a device (for example, fans and temperature sensors) and see related information such as models and serial numbers. You can also use the index numbers for these components to determine their operational status in the ENTITY-SENSOR-MIB and ENTITY-STATE-MIB.

Palo Alto Networks devices support only portions of the entPhysicalTable group:

| Object | Description |
|---|---|
| entPhysicalIndex | A single namespace that includes disk slots and disk drives. |
| entPhysicalDescr | The component description. |

| Object | Description |
|---|---|
| entPhysicalVendorType | The sysObjectID (see PAN-PRODUCT-MIB.my) when it is available (chassis and module objects). |
| entPhysicalContainedIn | The value of entPhysicalIndex for the component that contains this component. |
| entPhysicalClass | Chassis (3), container (5) for a slot, power supply (6), fan (7), sensor (8) for each temperature or other environmental, and module (9) for each line card. |
| entPhysicalParentRelPos | The relative position of this *child* component among its *sibling* components. Sibling components are defined as entPhysicalEntry components that share the same instance values of each of the entPhysicalContainedIn and entPhysicalClass objects. |
| entPhysicalName | Supported only if the management (MGT) interface allows for naming the line card. |
| entPhysicalHardwareRev | The vendor-specific hardware revision of the component. |
| entPhysicalFirwareRev | The vendor-specific firmware revision of the component. |
| entPhysicalSoftwareRev | The vendor-specific software revision of the component. |
| entPhysicalSerialNum | The vendor-specific serial number of the component. |
| entPhysicalMfgName | The name of the manufacturer of the component. |
| entPhysicalMfgDate | The date when the component was manufactured. |
| entPhysicalModelName | The disk model number. |
| entPhysicalAlias | An alias that the network manager specified for the component. |
| entPhysicalAssetID | A user-assigned asset tracking identifier that the network manager specified for the component. |
| entPhysicalIsFRU | Indicates whether the component is a field replaceable unit (FRU). |
| entPhysicalUris | The Common Language Equipment Identifier (CLEI) number of the component (for example, URN:CLEI:CNME120ARA). |

RFC 4133 defines this MIB.

## ENTITY-SENSOR-MIB

ENTITY-SENSOR-MIB adds support for physical sensors of networking equipment beyond what ENTITY-MIB defines. Use this MIB in tandem with the ENTITY-MIB to monitor the operational status of the physical components of a device (for example, fans and temperature sensors). For example, to troubleshoot issues that might result from environmental conditions, you can map the entity indexes from the ENTITY-MIB (entPhysicalDescr object) to operational status values (entPhysSensorOperStatus object) in the ENTITY-SENSOR-MIB. In the following example, all the fans and temperature sensors for a PA-3020 firewall are working:

 © Palo Alto Networks, Inc.

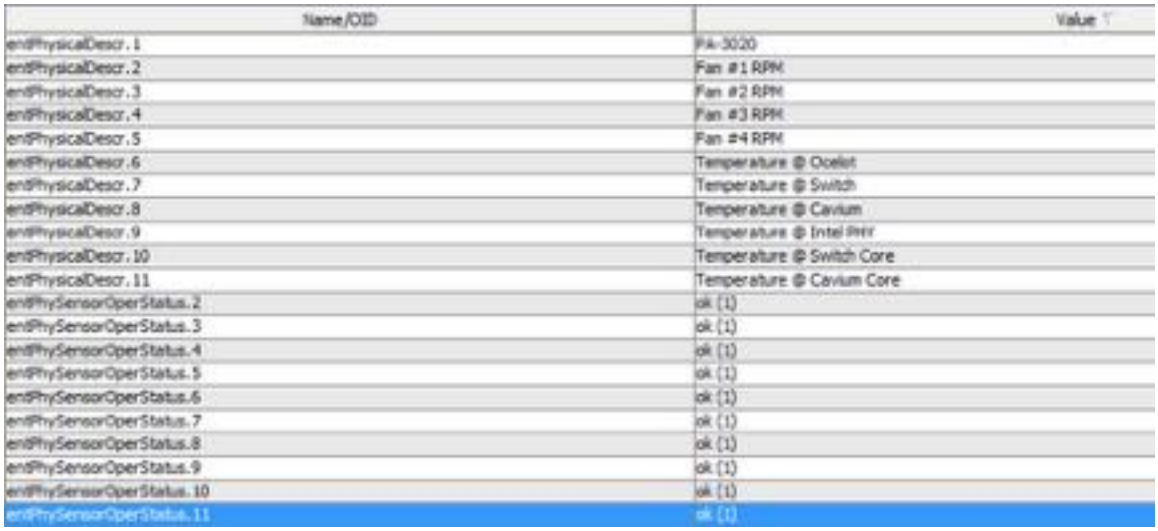

| Name/OID | Value |
|---|---|
| entPhysicalDescr.1 | PA-3020 |
| entPhysicalDescr.2 | Fan #1 RPM |
| entPhysicalDescr.3 | Fan #2 RPM |
| entPhysicalDescr.4 | Fan #3 RPM |
| entPhysicalDescr.5 | Fan #4 RPM |
| entPhysicalDescr.6 | Temperature @ Ocelot |
| entPhysicalDescr.7 | Temperature @ Switch |
| entPhysicalDescr.8 | Temperature @ Cavium |
| entPhysicalDescr.9 | Temperature @ Intel PHY |
| entPhysicalDescr.10 | Temperature @ Switch Core |
| entPhysicalDescr.11 | Temperature @ Cavium Core |
| entPhySensorOperStatus.2 | ok (1) |
| entPhySensorOperStatus.3 | ok (1) |
| entPhySensorOperStatus.4 | ok (1) |
| entPhySensorOperStatus.5 | ok (1) |
| entPhySensorOperStatus.6 | ok (1) |
| entPhySensorOperStatus.7 | ok (1) |
| entPhySensorOperStatus.8 | ok (1) |
| entPhySensorOperStatus.9 | ok (1) |
| entPhySensorOperStatus.10 | ok (1) |
| entPhySensorOperStatus.11 | ok (1) |

The same OID might refer to different sensors on different device platforms. Use the ENTITY-MIB for the targeted platform to match the value to the description.

Palo Alto Networks devices support only portions of the entPhySensorTable group. The supported portions vary by platform. The devices support only thermal (temperature in Celsius) and fan (in RPM) sensors.

RFC 3433 defines the ENTITY-SENSOR-MIB.

## ENTITY-STATE-MIB

ENTITY-STATE-MIB provides information about the state of physical components beyond what ENTITY-MIB defines, including the administrative and operational state of components in chassis-based platforms. Use this MIB in tandem with the ENTITY-MIB to monitor the operational state of the components of a PA-7000 Series firewall (for example, line cards, fan trays, and power supplies). For example, to troubleshoot log forwarding issues for Threat logs, you can map the log processing card (LPC) indexes from the ENTITY-MIB (entPhysicalDescr object) to operational state values (entStateOper object) in the ENTITY-STATE-MIB. The operational state values use numbers to indicate state: 1 for unknown, 2 for disabled, 3 for enabled, and 4 for testing. The PA-7000 Series firewall is the only Palo Alto Networks device that supports this MIB.

RFC 4268 defines the ENTITY-STATE-MIB.

## IEEE 802.3 LAG MIB

Use the IEEE 802.3 LAG MIB to monitor the status of aggregate groups that have Link Aggregation Control Protocol (LACP) enabled. When the firewall logs LACP events, it also generates traps that are useful for troubleshooting. For example, the traps can tell you whether traffic interruptions between the firewall and an LACP peer resulted from lost connectivity or from mismatched interface speed and duplex values.

PAN-OS implements the following SNMP tables for LACP. Note that the dot3adTablesLastChanged object indicates the time of the most recent change to dot3adAggTable, dot3adAggPortListTable, and dot3adAggPortTable.

| Table | Description |
|---|---|
| Aggregator Configuration Table (dot3adAggTable) | This table contains information about every aggregate group that is associated with a firewall. Each aggregate group has one entry.<br>Some table objects have restrictions, which the dot3adAggIndex object describes. This index is the unique identifier that the local system assigns to the aggregate group. It identifies an aggregate group instance among the subordinate managed objects of the containing object. The identifier is read-only.<br> The ifTable MIB (a list of interface entries) does not support logical interfaces and therefore does not have an entry for the aggregate group. |
| Aggregation Port List Table (dot3adAggPortListTable) | This table lists the ports associated with each aggregate group in a firewall. Each aggregate group has one entry.<br>The dot3adAggPortListPorts attribute lists the complete set of ports associated with an aggregate group. Each bit set in the list represents a port member. For non-chassis platforms, this is a 64-bit value. For chassis platforms, the value is an array of eight 64-bit entries. |
| Aggregation Port Table (dot3adAggPortTable) | This table contains LACP configuration information about every port associated with an aggregate group in a firewall. Each port has one entry. The table has no entries for ports that are not associated with an aggregate group. |
| LACP Statistics Table (dot3adAggPortStatsTable) | This table contains link aggregation information about every port associated with an aggregate group in a firewall. Each port has one row. The table has no entries for ports that are not associated with an aggregate group. |

The IEEE 802.3 LAG MIB includes the following LACP-related traps:

| Trap Name | Description |
|---|---|
| panLACPLostConnectivityTrap | The peer lost connectivity to the firewall. |
| panLACPUnresponsiveTrap | The peer does not respond to the firewall. |
| panLACPNegoFailTrap | LACP negotiation with the peer failed. |
| panLACPSpeedDuplexTrap | The link speed and duplex settings on the firewall and peer do not match. |
| panLACPLinkDownTrap | An interface in the aggregate group is down. |
| panLACPLacpDownTrap | An interface was removed from the aggregate group. |
| panLACPLacpUpTrap | An interface was added to the aggregate group. |

For the MIB definitions, refer to IEEE 802.3 LAG MIB.

 © Palo Alto Networks, Inc.

## LLDP-V2-MIB.my

Use the LLDP-V2-MIB to monitor Link Layer Discovery Protocol (LLDP) events. For example, you can check the lldpV2StatsRxPortFramesDiscardedTotal object to see the number of LLDP frames that were discarded for any reason. The Palo Alto Networks firewall uses LLDP to discover neighboring devices and their capabilities. LLDP makes troubleshooting easier, especially for virtual wire deployments where the ping or traceroute utilities won't detect the firewall.

Palo Alto Networks devices support all the LLDP-V2-MIB objects except:

- The following lldpV2Statistics objects:
  - lldpV2StatsRemTablesLastChangeTime
  - lldpV2StatsRemTablesInserts
  - lldpV2StatsRemTablesDeletes
  - lldpV2StatsRemTablesDrops
  - lldpV2StatsRemTablesAgeouts
- The following lldpV2RemoteSystemsData objects:
  - The lldpV2RemOrgDefInfoTable table
  - In the lldpV2RemTable table: lldpV2RemTimeMark

RFC 4957 defines this MIB.

## PAN-COMMON-MIB.my

Use the PAN-COMMON-MIB to monitor the following information for Palo Alto Networks devices:

| Object Group | Description |
|---|---|
| panSys | Contains such objects as device software/hardware versions, dynamic content versions, serial number, HA mode/state, and global counters.<br>The global counters include those related to Denial of Service (DoS), IP fragmentation, TCP state, and dropped packets. Tracking these counters enables you to monitor traffic irregularities that result from DoS attacks, device or connection faults, or resource limitations. PAN-COMMON-MIB supports global counters for firewalls but not for Panorama. |
| panChassis | Chassis type and M-Series appliance mode (Panorama or Log Collector). |
| panSession | Session utilization information. For example, the total number of active sessions on the firewall or a specific virtual system. |
| panMgmt | Status of the connection from the firewall to the Panorama management server. |
| panGlobalProtect | GlobalProtect gateway utilization as a percentage, maximum tunnels allowed, and number of active tunnels. |

| Object Group | Description |
|---|---|
| panLogCollector | Log Collector information such as the logging rate, log database storage duration (in days), and RAID disk usage. |

## PAN-GLOBAL-REG-MIB.my

PAN-GLOBAL-REG-MIB.my contains global, top-level OID definitions for various sub-trees of Palo Alto Networks enterprise MIB modules. This MIB doesn't contain objects for you to monitor; it is required only for referencing by other MIBs.

## PAN-GLOBAL-TC-MIB.my

PAN-GLOBAL-TC-MIB.my defines conventions (for example, character length and allowed characters) for the text values of objects in Palo Alto Networks enterprise MIB modules. All Palo Alto Networks products use these conventions. This MIB doesn't contain objects for you to monitor; it is required only for referencing by other MIBs.

## PAN-LC-MIB.my

PAN-LC-MIB.my contains definitions of managed objects that Log Collectors (M-Series appliances in Log Collector mode) implement. Use this MIB to monitor the logging rate, log database storage duration (in days), and disk usage (in MB) of each logical disk (up to four) on a Log Collector. For example, you can use this information to determine whether you should add more Log Collectors or forward logs to an external server (for example, a syslog server) for archiving.

## PAN-PRODUCT-MIB.my

PAN-PRODUCT-MIB.my defines sysObjectID OIDs for all Palo Alto Networks products. This MIB doesn't contain objects for you to monitor; it is required only for referencing by other MIBs.

## PAN-ENTITY-EXT-MIB.my

Use PAN-ENTITY-EXT-MIB.my in tandem with the ENTITY-MIB to monitor power usage for the physical components of a PA-7000 Series firewall (for example, fan trays, and power supplies), which is the only Palo Alto Networks device that supports this MIB. For example, when troubleshooting log forwarding issues, you might want to check the power usage of the log processing cards (LPCs): you can map the LPC indexes from the ENTITY-MIB (entPhysicalDescr object) to values in the PAN-ENTITY-EXT-MIB (panEntryFRUModelPowerUsed object).

 © Palo Alto Networks, Inc.

## PAN-TRAPS.my

Use PAN-TRAPS.my to see a complete listing of all the generated traps and information about them (for example, a description). For a list of traps that Palo Alto Networks devices support, refer to the PAN-COMMON-MIB.my **> panCommonEvents > panCommonEventsEvents > panCommonEventEventsV2** object.

# NetFlow Monitoring

NetFlow is an industry-standard protocol that the firewall can use to export statistics about the IP traffic that traverses its interfaces. The firewall exports the statistics as NetFlow fields to a NetFlow collector. The NetFlow collector is a server you use to analyze network traffic for security, administration, accounting and troubleshooting. All Palo Alto Networks firewalls support NetFlow (Version 9) except the PA-4000 Series and PA-7000 Series firewalls. The firewalls support only unidirectional NetFlow, not bidirectional. You can enable NetFlow exports on all interface types except HA, log card, or decrypt mirror. To identify firewall interfaces in a NetFlow collector, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors. The firewall supports standard and enterprise (PAN-OS specific) NetFlow templates.


- Configure NetFlow Exports
- NetFlow Templates


## Configure NetFlow Exports

| Configure NetFlow Exports | | |
|---|---|---|
| Step 1 | Create a NetFlow server profile. | 1. Select **Device > Server Profiles > NetFlow** and click **Add**.<br>2. Enter a **Name** for the profile.<br>3. Specify the frequency at which the firewall refreshes NetFlow Templates in **Minutes** (default is 30) or **Packets** (default is 20), according to the requirements of your NetFlow collector.<br>4. For the **Active Timeout**, specify the frequency in minutes at which the firewall exports records (default is 5).<br>5. Select the **PAN-OS Field Types** check box if you want the firewall to export App-ID and User-ID fields.<br>6. For each NetFlow collector (up to two per profile) that will receive fields, click **Add** and enter an identifying server **Name**, hostname or IP address (**NetFlow Server**), and access **Port** (default is 2055).<br>7. Click **OK** to save the profile. |
| Step 2 | Assign the NetFlow server profile to the interfaces that carry the traffic you want to analyze.<br>In this example, you assign the profile to an existing Ethernet interface. | 1. Select **Network > Interfaces > Ethernet** and click an interface name to edit it.<br>2. In the **NetFlow Profile** drop-down, select the NetFlow server profile and click **OK**.<br>3. Click **Commit**. |
| Step 3 | Monitor the firewall traffic in a NetFlow collector. | Refer to the documentation for your NetFlow collector.<br>When monitoring statistics, you must match the interface indexes in the NetFlow collector with interface names in the firewall web interface. For details, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors. |

 © Palo Alto Networks, Inc.

## NetFlow Templates

NetFlow collectors use templates to decipher the fields that the firewall exports. The firewall selects a template based on the type of exported data: IPv4 or IPv6 traffic, with or without NAT, and with standard or enterprise-specific (PAN-OS specific) fields. The firewall periodically refreshes templates to re-evaluate which one to use (in case the type of exported data changes) and to apply any changes to the fields in the selected template. When you Configure NetFlow Exports, you set the refresh frequency according to the requirements of your NetFlow collector.

The Palo Alto Networks firewall supports the following NetFlow templates:

| Template | ID |
|---|---|
| IPv4 Standard | 256 |
| IPv4 Enterprise | 257 |
| IPv6 Standard | 258 |
| IPv6 Enterprise | 259 |
| IPv4 with NAT Standard | 260 |
| IPv4 with NAT Enterprise | 261 |
| IPv6 with NAT Standard | 262 |
| IPv6 with NAT Enterprise | 263 |

The following table lists the NetFlow fields that the firewall can send, along with the templates that define them:

| Value | Field | Description | Templates |
|---|---|---|---|
| 1 | IN_BYTES | Incoming counter with length N * 8 bits for the number of bytes associated with an IP flow. By default, N is 4. | All templates |
| 2 | IN_PKTS | Incoming counter with length N * 8 bits for the number of packets associated with an IP glow. By default, N is 4. | All templates |
| 4 | PROTOCOL | IP protocol byte. | All templates |
| 5 | TOS | Type of Service byte setting when entering the ingress interface. | All templates |
| 6 | TCP_FLAGS | Total of all the TCP flags in this flow. | All templates |
| 7 | L4_SRC_PORT | TCP/UDP source port number (for example, FTP, Telnet, or equivalent). | All templates |
| 8 | IPV4_SRC_ADDR | IPv4 source address. | IPv4 standard<br>IPv4 enterprise<br>IPv4 with NAT standard<br>IPv4 with NAT enterprise |

| Value | Field | Description | Templates |
|---|---|---|---|
| 10 | INPUT_SNMP | Input interface index. The value length is 2 bytes by default, but higher values are possible. For details on how Palo Alto Networks firewalls generate interface indexes, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors. | All templates |
| 11 | L4_DST_PORT | TCP/UDP destination port number (for example, FTP, Telnet, or equivalent). | All templates |
| 12 | IPV4_DST_ADDR | IPv4 destination address. | IPv4 standard<br>IPv4 enterprise<br>IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 14 | OUTPUT_SNMP | Output interface index. The value length is 2 bytes by default, but higher values are possible. For details on how Palo Alto Networks firewalls generate interface indexes, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors. | All templates |
| 21 | LAST_SWITCHED | System uptime in milliseconds when the last packet of this flow was switched. | All templates |
| 22 | FIRST_SWITCHED | System uptime in milliseconds when the first packet of this flow was switched. | All templates |
| 27 | IPV6_SRC_ADDR | IPv6 source address. | IPv6 standard<br>IPv6 enterprise<br>IPv6 with NAT standard<br>IPv6 with NAT enterprise |
| 28 | IPV6_DST_ADDR | IPv6 destination address. | IPv6 standard<br>IPv6 enterprise<br>IPv6 with NAT standard<br>IPv6 with NAT enterprise |
| 32 | ICMP_TYPE | Internet Control Message Protocol (ICMP) packet type. This is reported as:<br>ICMP Type * 256 + ICMP code | All templates |
| 61 | DIRECTION | Flow direction:<br>• 0 = ingress<br>• 1 = egress | All templates |

 © Palo Alto Networks, Inc.

| Value | Field | Description | Templates |
|---|---|---|---|
| 148 | flowId | An identifier of a flow that is unique within an observation domain. You can use this information element to distinguish between different flows if flow keys such as IP addresses and port numbers are not reported or are reported in separate records. The flowID corresponds to the session ID field in Traffic and Threat logs. | All templates |
| 233 | firewallEvent | Indicates a firewall event:<br>• 0 = Ignore (invalid)<br>• 1 = Flow created<br>• 2 = Flow deleted<br>• 3 = Flow denied<br>• 4 = Flow alert<br>• 5 = Flow update (the session state changed from active to deny) | All templates |
| 225 | postNATSourceIPv4Address | The definition of this information element is identical to that of sourceIPv4Address, except that it reports a modified value that the firewall produced during network address translation after the packet traversed the interface. | IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 226 | postNATDestinationIPv4Address | The definition of this information element is identical to that of destinationIPv4Address, except that it reports a modified value that the firewall produced during network address translation after the packet traversed the interface. | IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 227 | postNAPTSourceTransportPort | The definition of this information element is identical to that of sourceTransportPort, except that it reports a modified value that the firewall produced during network address port translation after the packet traversed the interface. | IPv4 with NAT standard<br>IPv4 with NAT enterprise |
| 228 | postNAPTDestinationTransportPort | The definition of this information element is identical to that of destinationTransportPort, except that it reports a modified value that the firewall produced during network address port translation after the packet traversed the interface. | IPv4 with NAT standard<br>IPv4 with NAT enterprise |

© Palo Alto Networks, Inc.

| Value | Field | Description | Templates |
|---|---|---|---|
| 281 | postNATSourceIPv6Address | The definition of this information element is identical to the definition of information element sourceIPv6Address, except that it reports a modified value that the firewall produced during NAT64 network address translation after the packet traversed the interface. See RFC 2460 for the definition of the source address field in the IPv6 header. See RFC 6146 for NAT64 specification. | IPv6 with NAT standard<br>IPv6 with NAT enterprise |
| 282 | postNATDestinationIPv6Address | The definition of this information element is identical to the definition of information element destinationIPv6Address, except that it reports a modified value that the firewall produced during NAT64 network address translation after the packet traversed the interface. See RFC 2460 for the definition of the destination address field in the IPv6 header. See RFC 6146 for NAT64 specification. | IPv6 with NAT standard<br>IPv6 with NAT enterprise |
| 346 | privateEnterpriseNumber | This is a unique private enterprise number that identifies Palo Alto Networks: 25461. | IPv4 enterprise<br>IPv4 with NAT enterprise<br>IPv6 enterprise<br>IPv6 with NAT enterprise |
| 56701 | App-ID | The name of an application that App-ID identified. The name can be up to 32 bytes. | IPv4 enterprise<br>IPv4 with NAT enterprise<br>IPv6 enterprise<br>IPv6 with NAT enterprise |
| 56702 | User-ID | A username that User-ID identified. The name can be up to 64 bytes. | IPv4 enterprise<br>IPv4 with NAT enterprise<br>IPv6 enterprise<br>IPv6 with NAT enterprise |

 © Palo Alto Networks, Inc.

# Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors

When you use a NetFlow collector (see NetFlow Monitoring) or SNMP manager (see SNMP Monitoring and Traps) to monitor the Palo Alto Networks firewall, an interface index (SNMP ifindex object) identifies the interface that carried a particular flow (see Figure: Interface Indexes in an SNMP Manager). In contrast, the firewall web interface uses interface names as identifiers (for example, ethernet1/1), not indexes. To understand which statistics that you see in a NetFlow collector or SNMP manager apply to which firewall interface, you must be able to match the interface indexes with interface names.

**Figure: Interface Indexes in an SNMP Manager**

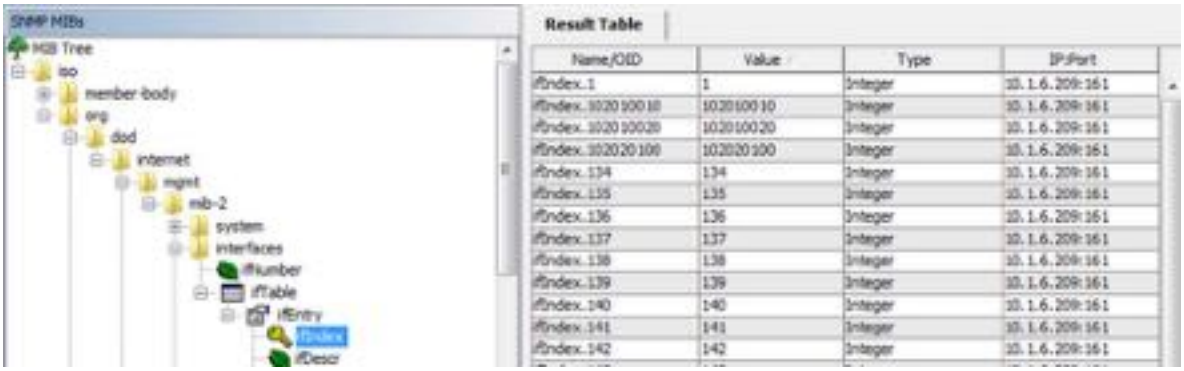


You can match the indexes with names by understanding the formulas that the firewall uses to calculate indexes. The formulas vary by platform and interface type: physical or logical.

Physical interface indexes have a range of 1-9999, which the firewall calculates as follows:

| Firewall Platform | Calculation | Example Interface Index |
|---|---|---|
| Non-chassis based:<br>VM-Series, PA-200, PA-500, PA-2000 Series, PA-3000 Series, PA-4000 Series, PA-5000 Series<br>The PA-4000 Series platform supports SNMP but not NetFlow. | MGT port + physical port offset<br>• MGT port—This is a constant that depends on the platform:<br>  • 2 for hardware-based firewalls (for example, the PA-5000 Series firewall)<br>  • 1 for the VM-Series firewall<br>• Physical port offset—This is the physical port number. | PA-5000 Series firewall, Eth1/4 =<br>2 (MGT port) + 4 (physical port) = **6** |

| Firewall Platform | Calculation | Example Interface Index |
|---|---|---|
| Chassis based:<br>PA-7000 Series firewalls<br>This platform supports SNMP but not NetFlow. | (Max. ports * slot) + physical port offset + MGT port<br>• Maximum ports—This is a constant of 64.<br>• Slot—This is the chassis slot number of the network interface card.<br>• Physical port offset—This is the physical port number.<br>• MGT port—This is a constant of 5 for PA-7000 Series firewalls. | PA-7000 Series firewall, Eth3/9 =<br>[64 (max. ports) * 3 (slot)] + 9 (physical port) + 5 (MGT port) = **206** |

Logical interface indexes for all platforms are nine-digit numbers that the firewall calculates as follows:

| Interface Type | Range | Digit 9 | Digits 7-8 | Digits 5-6 | Digits 1-4 | Example Interface Index |
|---|---|---|---|---|---|---|
| Layer 3 subinterface | 101010001-199999999 | Type: 1 | Interface slot: 1-9 (01-09) | Interface port: 1-9 (01-09) | Subinterface: suffix 1-9999 (0001-9999) | Eth1/5.22 = 100000000 (type) + 100000 (slot) + 50000 (port) + 22 (suffix) = **101050022** |
| Layer 2 subinterface | 101010001-199999999 | Type: 1 | Interface slot: 1-9 (01-09) | Interface port: 1-9 (01-09) | Subinterface: suffix 1-9999 (0001-9999) | Eth2/3.6 = 100000000 (type) + 200000 (slot) + 30000 (port) + 6 (suffix) = **102030006** |
| Vwire subinterface | 101010001-199999999 | Type: 1 | Interface slot: 1-9 (01-09) | Interface port: 1-9 (01-09) | Subinterface: suffix 1-9999 (0001-9999) | Eth4/2.312 = 100000000 (type) + 400000 (slot) + 20000 (port) + 312 (suffix) = **104020312** |
| VLAN | 200000001-200009999 | Type: 2 | 00 | 00 | VLAN suffix: 1-9999 (0001-9999) | VLAN.55 = 200000000 (type) + 55 (suffix) = **200000055** |
| Loopback | 300000001-300009999 | Type: 3 | 00 | 00 | Loopback suffix: 1-9999 (0001-9999) | Loopback.55 = 300000000 (type) + 55 (suffix) = **300000055** |
| Tunnel | 400000001-400009999 | Type: 4 | 00 | 00 | Tunnel suffix: 1-9999 (0001-9999) | Tunnel.55 = 400000000 (type) + 55 (suffix) = **400000055** |
| Aggregate group | 500010001-500089999 | Type: 5 | 00 | AE suffix: 1-8 (01-08) | Subinterface: suffix 1-9999 (0001-9999) | AE5.99 = 500000000 (type) + 50000 (AE Suffix) + 99 (suffix) = **500050099** |


© Palo Alto Networks, Inc.

# User-ID

The User Identification (User-ID™) feature of the Palo Alto Networks next-generation firewall enables you to create policies and perform reporting based on users and groups rather than individual IP addresses.


- User-ID Overview
- User-ID Concepts
- Enable User-ID
- Map Users to Groups
- Map IP Addresses to Users
- Configure a Firewall to Share User Mapping Data with Other Firewalls
- Enable User- and Group-Based Policy
- Enable Policy for Users with Multiple Accounts
- Verify the User-ID Configuration
- Deploy User-ID in a Large-Scale Network

# User-ID Overview

User-ID seamlessly integrates Palo Alto Networks firewalls with a range of enterprise directory and terminal services offerings, enabling you to tie application activity and policy rules to users and groups—not just IP addresses. Furthermore, with User-ID enabled, the Application Command Center (ACC), App-Scope, reports, and logs all include usernames in addition to user IP addresses.

Palo Alto Networks firewalls support monitoring of the following enterprise services:

- Microsoft Active Directory
- Lightweight Directory Access Protocol (LDAP)
- Novell eDirectory
- Citrix Metaframe Presentation Server or XenApp
- Microsoft Terminal Services

For user- and group-based policies, the firewall requires a list of all available users and their corresponding group mappings that you can select when defining your policies. The firewall collects Group Mapping information by connecting directly to your LDAP directory server.

To enforce user- and group-based policies, the firewall must be able to map the IP addresses in the packets it receives to usernames. User-ID provides many mechanisms to collect this User Mapping information. For example, the User-ID agent monitors server logs for login events, probes clients, and listens for syslog messages from authenticating services. To identify mappings for IP addresses that the agent didn't map, you can configure the firewall to redirect HTTP requests to a Captive Portal login. You can tailor the user mapping mechanisms to suit your environment, and even use different mechanisms at different sites.

User-ID does not work in environments where the source IP addresses of users are subject to NAT translation before the firewall maps the IP addresses to usernames.

 © Palo Alto Networks, Inc.

**Figure: User-ID**

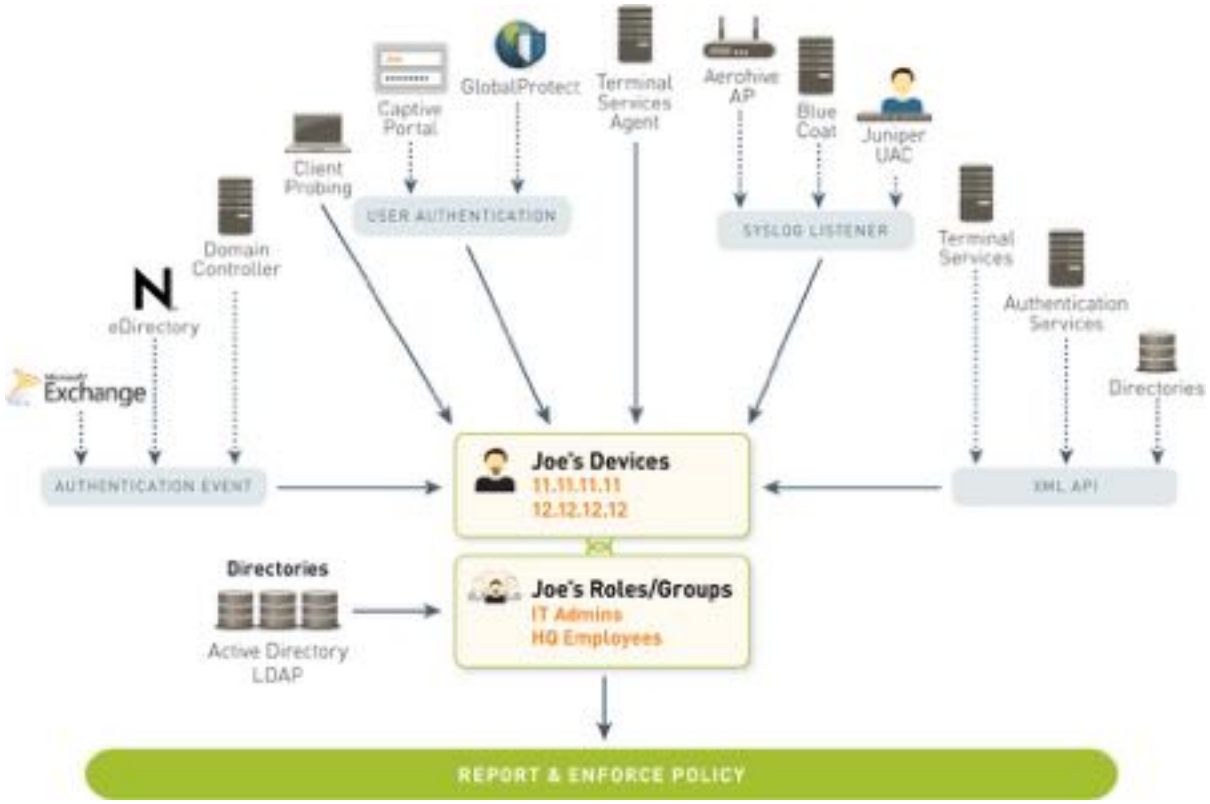


See User-ID Concepts for information on how User-ID works and Enable User-ID for instructions on setting up User-ID.

# User-ID Concepts

▲ Group Mapping
▲ User Mapping

## Group Mapping

To define policy rules based on user or group, first you create an LDAP server profile that defines how the firewall connects and authenticates to your directory server. The firewall supports a variety of directory servers, including Microsoft Active Directory (AD), Novell eDirectory, and Sun ONE Directory Server. The server profile also defines how the firewall searches the directory to retrieve the list of groups and the corresponding list of members. Next you create a group mapping configuration to Map Users to Groups. Then you can select the users or groups when defining policy rules.

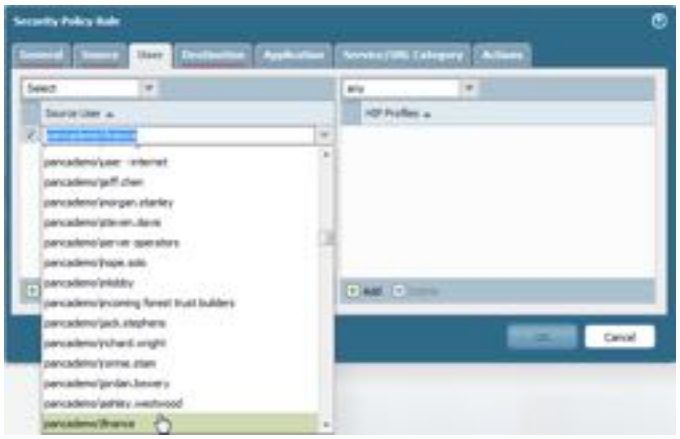

Defining policy rules based on group membership rather than on individual users simplifies administration because you don't have to update the rules whenever new users are added to a group. For example, the following security rules allow access to specific internal applications based on group membership:

| Name | Tags | Type | Zone | Address | User | HIP Profile | Application |
|---|---|---|---|---|---|---|---|
| CRM Access | none | universal | trust | any | Finance | Patches | sap |
| Eng access | none | universal | trust | any | Engineering | any | bugzilla |
| | | | | | | | perforce |

When configuring group mapping, you can limit which groups will be available in policy rules. You can specify groups that already exist in your directory service or define custom groups based on LDAP filters. Defining custom groups can be quicker than creating new groups or changing existing ones on an LDAP server, and doesn't require an LDAP administrator to intervene. User-ID maps all the LDAP directory users who match the filter to the custom group. For example, you might want a security policy that allows contractors in the Marketing Department to access social networking sites. If no Active Directory group exists for that department, you can configure an LDAP filter that matches users for whom the LDAP attribute Department is set to Marketing. Log queries and reports that are based on user groups will include custom groups.

 © Palo Alto Networks, Inc.

# User Mapping

Having the names of the users and groups is only one piece of the puzzle. The firewall also needs to know which IP addresses map to which users so that security rules can be enforced appropriately. Figure: User-ID illustrates the different methods that are used to identify users and groups on your network and shows how user mapping and group mapping work together to enable user- and group-based security enforcement and visibility.

The following topics describe the different methods of user mapping:

- Server Monitoring
- Client Probing
- Port Mapping
- Syslog
- Captive Portal
- GlobalProtect
- User-ID XML API

## Server Monitoring

With server monitoring a User-ID agent—either a Windows-based agent running on a domain server in your network, or the integrated PAN-OS User-ID agent running on the firewall—monitors the security event logs for specified Microsoft Exchange Servers, domain controllers, or Novell eDirectory servers for login events. For example, in an AD environment, you can configure the User-ID agent to monitor the security logs for Kerberos ticket grants or renewals, Exchange server access (if configured), and file and print service connections. Note that for these events to be recorded in the security log, the AD domain must be configured to log successful account login events. In addition, because users can log in to any of the servers in the domain, you must set up server monitoring for all servers to capture all user login events.

Because server monitoring requires very little overhead and because the majority of users can generally be mapped using this method, it is recommended as the base user mapping method for most User-ID deployments. See Configure User Mapping Using the Windows User-ID Agent or Configure User Mapping Using the PAN-OS Integrated User-ID Agent for details.

## Client Probing

In a Microsoft Windows environment, you can configure the User-ID agent to probe client systems using Windows Management Instrumentation (WMI). The Windows-based User-ID agent can also perform NetBIOS probing (not supported on the PAN-OS integrated User-ID agent). Probing is particularly useful in environments with a high IP address turnover because changes will be reflected on the firewall more quickly, enabling more accurate enforcement of user-based policies. However, if the correlation between IP addresses and users is fairly static, you probably do not need to enable client probing. Because probing can generate a large amount of network traffic (based on the total number of mapped IP addresses), the agent that will be initiating the probes should be located as close as possible to the end clients.

© Palo Alto Networks, Inc.

If probing is enabled, the agent will probe each learned IP address periodically (every 20 minutes by default, but this is configurable) to verify that the same user is still logged in. In addition, when the firewall encounters an IP address for which it has no user mapping, it will send the address to the agent for an immediate probe.

See Configure User Mapping Using the Windows User-ID Agent or Configure User Mapping Using the PAN-OS Integrated User-ID Agent for details.

## Port Mapping

In environments with multi-user systems—such as Microsoft Terminal Server or Citrix environments—many users share the same IP address. In this case, the user-to-IP address mapping process requires knowledge of the source port of each client. To perform this type of mapping, you must install the Palo Alto Networks Terminal Services Agent on the Windows/Citrix terminal server itself to intermediate the assignment of source ports to the various user processes. For terminal servers that do not support the Terminal Services Agent, such as Linux terminal servers, you can use the XML API to send user mapping information from login and logout events to User-ID. See Configure User Mapping for Terminal Server Users for configuration details.

## Syslog

In environments with existing network services that authenticate users—such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms—the firewall User-ID agent (either the Windows agent or the PAN-OS integrated agent on the firewall) can listen for authentication syslog messages from those services. Syslog filters, which are provided by a content update (integrated User-ID agent only) or configured manually, allow the User-ID agent to parse and extract usernames and IP addresses from authentication syslog events generated by the external service, and add the information to the User-ID IP address-to-username mappings maintained by the firewall. See Configure User-ID to Receive User Mappings from a Syslog Sender for configuration details.

 © Palo Alto Networks, Inc.

**Figure: User-ID Integration with Syslog**

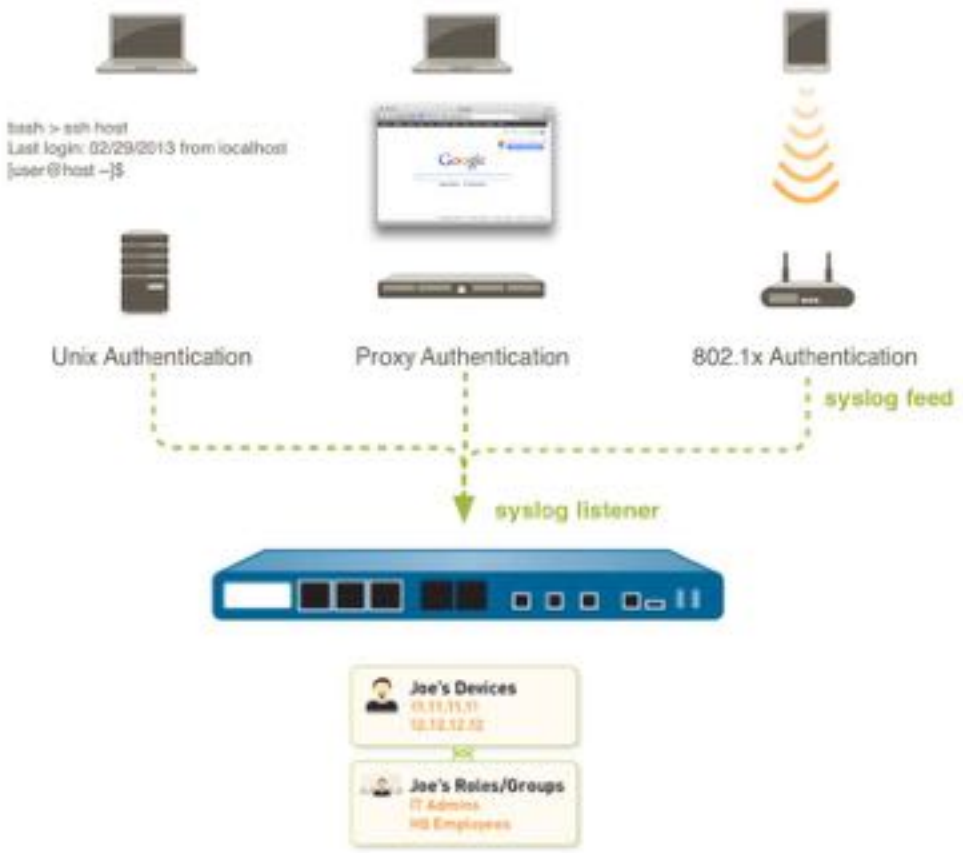


## Captive Portal

If the firewall or the User-ID agent can't map an IP address to a username—for example, if the user isn't logged in or uses an operating system such as Linux that your domain servers don't support—you can configure Captive Portal. Any web traffic (HTTP or HTTPS) that matches a Captive Portal policy rule requires user authentication. You can base the authentication on a transparent browser-challenge (Kerberos Single Sign-On (SSO) or NT LAN Manager (NTLM) authentication), web form (for RADIUS, TACACS+, LDAP, Kerberos, or local database authentication), or client certificates. For details, see Map IP Addresses to Usernames Using Captive Portal.

## GlobalProtect

For mobile or roaming users, the GlobalProtect client provides the user mapping information to the firewall directly. In this case, every GlobalProtect user has an agent or app running on the client that requires the user to enter login credentials for VPN access to the firewall. This login information is then added to the User-ID user mapping table on the firewall for visibility and user-based security policy enforcement. Because GlobalProtect users must authenticate to gain access to the network, the IP address-to-username mapping is

explicitly known. This is the best solution in sensitive environments where you must be certain of who a user is in order to allow access to an application or service. For more information on setting up GlobalProtect, refer to the GlobalProtect Administrator's Guide.

## User-ID XML API

For other types of user access that cannot be mapped using any of the standard user mapping methods or Captive Portal—for example, to add mappings of users connecting from a third-party VPN solution or users connecting to a 802.1x enabled wireless network—you can use the User-ID XML API to capture login events and send them to the User-ID agent or directly to the firewall. See Send User Mappings to User-ID Using the XML API for details.

 © Palo Alto Networks, Inc.

# Enable User-ID

You must complete the following tasks to set up the firewall to user users and groups in policy enforcement, logging, and reporting:

- Map Users to Groups
- Map IP Addresses to Users
- Enable User- and Group-Based Policy
- Verify the User-ID Configuration

© Palo Alto Networks, Inc.

# Map Users to Groups

Use the following procedure to enable the firewall to connect to your LDAP directory and retrieve Group Mapping information:

The following are best practices for group mapping in an Active Directory (AD) environment:

- If you have a single domain, you need only one LDAP server profile that connects the firewall to the domain controller with the best connectivity. You can add additional domain controllers for fault tolerance.
- If you have multiple domains and/or multiple forests, you must create a server profile to connect to a domain server in each domain/forest. Take steps to ensure unique usernames in separate forests.
- If you have Universal Groups, create a server profile to connect to the Global Catalog server.

| Map Users to Groups | |
|---|---|
| Step 1 Add an LDAP server profile.<br><br>The profile defines how the firewall connects to the directory servers from which it collects group mapping information. You can add up to four servers to the profile but they must be the same **Type**. | Configure an LDAP Server Profile:<br>1. Select **Device > Server Profiles > LDAP**, click **Add**, and enter a **Profile Name**.<br>2. For each LDAP server, click **Add** and enter the server **Name**, IP address (**LDAP Server**), and **Port** (default is 389).<br>3. Based on your **Type** selection (for example, **active-directory**), the firewall automatically populates the correct LDAP attributes in the group mapping settings. However, if you customized your LDAP schema, you might need to modify the default settings.<br>4. In the **Base DN** field, enter the Distinguished Name (DN) of the LDAP tree location where you want the firewall to begin its search for user and group information.<br>5. Enter the authentication credentials for binding to the LDAP tree in the **Bind DN**, **Password**, and **Confirm Password** fields. The **Bind DN** can be a fully qualified LDAP name (for example, `cn=administrator,cn=users,dc=acme,dc=local`) or a user principal name (for example, `administrator@acme.local`).<br>6. Click **OK** to save the profile. |


© Palo Alto Networks, Inc.

| Map Users to Groups (Continued) | |
|---|---|
| Step 2 Configure the server settings in a group mapping configuration. | 1. Select **Device > User Identification > Group Mapping Settings**.<br>2. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for this configuration.<br>3. Click **Add** and enter a unique **Name** to identify the group mapping configuration.<br>4. Select the LDAP **Server Profile** you just created.<br>5. (Optional) By default, the **User Domain** field is blank: the firewall automatically detects the domain names for Active Directory (AD) servers. If you enter a value, it overrides any domain names that the firewall retrieves from the LDAP source. Your entry must be the NetBIOS domain name.<br>6. (Optional) To filter the groups that the firewall tracks for group mapping, in the Group Objects section, enter a **Search Filter** (LDAP query), **Object Class** (group definition), **Group Name**, and **Group Member**.<br>7. (Optional) To filter the users that the firewall tracks for group mapping, in the User Objects section, enter a **Search Filter** (LDAP query), **Object Class** (user definition), and **User Name**.<br>8. (Optional) To match User-ID information with email header information identified in the links and attachments of emails forwarded to WildFire™, enter the list of email domains in your organization in the Mail Domains section, **Domain List** field. Use commas to separate multiple domains (up to 256 characters). After you click **OK**, PAN-OS automatically populates the **Mail Attributes** field based on your LDAP server type (Sun/RFC, Active Directory, or Novell). When a match occurs, the username in the WildFire log email header section will contain a link that opens the **ACC** tab, filtered by user or user group.<br>9. Make sure the **Enabled** check box is selected. |

© Palo Alto Networks, Inc.

| Map Users to Groups (Continued) | |
|---|---|
| Step 3 Limit which groups will be available in policy rules.<br><br>Required only if you want to limit policy rules to specific groups. By default, if you don't specify groups, all groups are available in policy rules.<br><br> Any custom groups you create will also be available in the Allow List of authentication profiles. | 1. Add existing groups from the directory service:<br>a. Select the **Group Include List** tab.<br>b. In the Available Groups list, select the groups you want to appear in policy rules and click the Add icon.<br>2. If you want to base policy rules on user attributes that don't match existing user groups, create custom groups based on LDAP filters:<br>a. Select the **Custom Group** tab and click **Add**.<br>b. Enter a group **Name** that is unique in the group mapping configuration for the current firewall or virtual system. If the **Name** has the same value as the Distinguished Name (DN) of an existing AD group domain, the firewall uses the custom group in all references to that name (for example, in policies and logs).<br>c. Specify an **LDAP Filter** of up to 2,048 UTF-8 characters and click **OK**. The firewall doesn't validate LDAP filters, so it's up to you to ensure they are accurate.<br><br> To minimize the performance impact on the LDAP directory server, use only indexed attributes in the filter.<br>3. Click **OK** and **Commit**. A commit is necessary before custom groups will be available in policies and objects. |

 © Palo Alto Networks, Inc.

# Map IP Addresses to Users

The tasks you perform to map IP addresses to usernames depends on the type and location of the client systems on your network. Complete as many of the following tasks as necessary to enable mapping of your client systems:

- To map users as they log in to your Exchange servers, domain controllers, or eDirectory servers, or Windows clients that can be directly probed, you must configure a User-ID agent to monitor the server logs and probe client systems. You can either Configure User Mapping Using the Windows User-ID Agent (a standalone agent that you install on one or more member servers in the domain that contains the servers and clients that the agent will monitor) or Configure User Mapping Using the PAN-OS Integrated User-ID Agent. For guidance on which agent is appropriate for your network and the required number and placements of agents, refer to Architecting User Identification Deployments.
- If you have clients running multi-user systems in a Windows environment, such as Microsoft Terminal Server or Citrix Metaframe Presentation Server or XenApp, Configure the Palo Alto Networks Terminal Server Agent for User Mapping. For a multi-user system that doesn't run on Windows, you can Retrieve User Mappings from a Terminal Server Using the User-ID XML API.
- To obtain user mappings from existing network services that authenticate users—such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms—Configure User-ID to Receive User Mappings from a Syslog Sender. You can use either the Windows agent or the agentless user mapping feature on the firewall to listen for authentication syslog messages from the network services.
- If you have users with client systems that aren't logged into your domain servers—for example, users running Linux clients that don't log in to the domain—you can Map IP Addresses to Usernames Using Captive Portal.
- For other clients that you can't map using the preceding methods, you can Send User Mappings to User-ID Using the XML API.
- Because policy is local to each firewall, each firewall needs current user mapping and group mapping information to accurately enforce policy by user and group. However, if you want one firewall to function as the sole, central collection and distribution point for user mappings, you can Configure a Firewall to Share User Mapping Data with Other Firewalls.

# Configure User Mapping Using the Windows User-ID Agent

In most cases, the majority of your network users will have logins to your monitored domain services. For these users, the Palo Alto Networks User-ID agent monitors the servers for login events and performs the IP address to username mapping. The way you configure the User-ID agent depends on the size of your environment and the location of your domain servers. As a best practice, you should locate your User-ID agents near your monitored servers (that is, the monitored servers and the Windows User-ID agent should not be across a WAN link from each other). This is because most of the traffic for user mapping occurs between the agent and the monitored server, with only a small amount of traffic—the delta of IP address mappings since the last update—from the agent to the firewall.

The following topics describe how to install and configure the User-ID Agent and how to configure the firewall to retrieve user mapping information from the agent:

- Install the User-ID Agent
- Configure the User-ID Agent for User Mapping

## Install the User-ID Agent

The following procedure shows how to install the User-ID agent on a member server in the domain and set up the service account with the required permissions. If you are upgrading, the installer will automatically remove the older version, however, it is a good idea to back up the config.xml file before running the installer.

For information about the system requirements for installing the Windows-based User-ID agent and for information on the supported server OS versions are supported, refer to "Operating System (OS) Compatibility User-ID Agent" in the User-ID Agent Release Notes, which are available on the Palo Alto Networks Software Updates page.

| Install the Windows User-ID Agent | | |
|---|---|---|
| Step 1 | Decide where to install the User-ID agent.<br><br>The User-ID agent queries the Domain Controller and Exchange server logs using Microsoft Remote Procedure Calls (MSRPCs), which require a complete transfer of the entire log at each query. Therefore, always install one or more User-ID agents at each site that has servers to be monitored.<br><br> For more detailed information on where to install User-ID agents, refer to Architecting User Identification (User-ID) Deployments. | • You must install the User-ID agent on a system running one of the supported OS versions: see "Operating System (OS) Compatibility User-ID Agent" in the User-ID Agent Release Notes.<br>• Make sure the system that will host the User-ID agent is a member of the same domain as the servers it will monitor.<br>• As a best practice, install the User-ID agent close to the servers it will be monitoring (there is more traffic between the User-ID agent and the monitored servers than there is between the User-ID agent and the firewall, so locating the agent close to the monitored servers optimizes bandwidth usage).<br>• To ensure the most comprehensive mapping of users, you must monitor all servers that contain user login information. You might need to install multiple User-ID agents to efficiently monitor all of your resources. |

 © Palo Alto Networks, Inc.

| Install the Windows User-ID Agent (Continued) | |
|---|---|
| Step 2 Download the User-ID agent installer.<br><br>As a best practice, install the User-ID agent version that is the same as the PAN-OS version running on the firewalls. | 1. Log in to Palo Alto Networks Support site.<br>2. Select **Software Updates** from the Manage Devices section.<br>3. Scroll to the User Identification Agent section of the screen and **Download** the version of the User-ID agent you want to install.<br>4. Save the `UaInstall-x.x.x-xx.msi` file on the system(s) where you plan to install the agent. |
| Step 3 Run the installer as an administrator.<br> | 1. To launch a command prompt as an administrator, click Start and right-click **Command Prompt** and then select **Run as administrator**.<br>2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop you would enter the following:<br>`C:\Users\administrator.acme>cd Desktop`<br>`C:\Users\administrator.acme\Desktop>UaInstall-6.0.0-1.msi`<br>3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to the `C:\Program Files (x86)\Palo Alto Networks\User-ID Agent` folder, but you can **Browse** to a different location.<br>4. When the installation completes, **Close** the setup window. |
| Step 4 Launch the User-ID Agent application.<br> 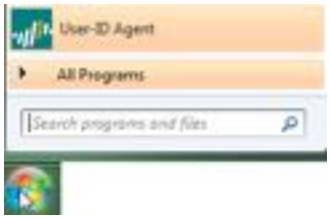 | 1. Click Start and select **User-ID Agent**. |
| Step 5 (Optional) Change the service account that the User-ID agent uses to log in.<br> | By default, the agent uses the administrator account used to install the .msi file. However, you may want to switch this to a restricted account as follows:<br>1. Select **User Identification > Setup** and click **Edit**.<br>2. Select the **Authentication** tab and enter the service account name that you want the User-ID agent to use in the **User name for Active Directory** field.<br>3. Enter the **Password** for the specified account. |

| Install the Windows User-ID Agent (Continued) | |
|---|---|
| Step 6 (Optional) Assign account permissions to the installation folder.<br><br>You only need to perform this step if the service account you configured for the User-ID agent is not a member of the administrators group for the domain or a member of both the Server Operators and the Event Log Readers groups. | 1. Give the service account permissions to the installation folder:<br>a. From the Windows Explorer, navigate to `C:\Program Files\Palo Alto Networks` and right-click the folder and select **Properties**.<br>b. On the **Security** tab, **Add** the User-ID agent service account and assign it permissions to **Modify**, **Read & execute**, **List folder contents**, and **Read** and then click **OK** to save the account settings.<br>2. Give the service account permissions to the User-ID Agent registry sub-tree:<br>a. Run `regedit32` and navigate to the Palo Alto Networks sub-tree in one of the following locations:<br>– **32-bit systems**—`HKEY_LOCAL_MACHINE\Software\Palo Alto Networks`<br>– **64-bit systems**—`HKEY_LOCAL_MACHINE\Software\WOW6432Node\Palo Alto Networks`<br>b. Right-click the Palo Alto Networks node and select **Permissions**.<br>c. Assign the User-ID service account **Full Control** and then click **OK** to save the setting.<br>3. On the domain controller, add the service account to the builtin groups to enable privileges to read the security log events (Event Log Reader group) and open sessions (Server Operator group):<br>a. Run the MMC and Launch the Active Directory Users and Computers snap-in.<br>b. Navigate to the Builtin folder for the domain and then right-click each group you need to edit (Event Log Reader and Server Operator) and select **Add to Group** to open the properties dialog.<br>c. Click **Add** and enter the name of the service account that you configured the User-ID service to use and then click **Check Names** to validate that you have the proper object name.<br>d. Click **OK** twice to save the settings. |

## Configure the User-ID Agent for User Mapping

The Palo Alto Networks User-ID agent is a Windows service that connects to servers on your network—for example, Active Directory servers, Microsoft Exchange servers, and Novell eDirectory servers—and monitors the logs for login events. The agent uses this information to map IP addresses to usernames. Palo Alto Networks firewalls connect to the User-ID agent to retrieve this user mapping information, enabling visibility into user activity by username rather than IP address and enables user- and group-based security enforcement.

 © Palo Alto Networks, Inc.

For information about the server OS versions supported by the User-ID agent, refer to "Operating System (OS) Compatibility User-ID Agent" in the *User-ID Agent Release Notes*, which are available on the Palo Alto Networks Software Updates page.

| Map IP Addresses to Users Using the User-ID Agent | | |
|---|---|---|
| Step 1 | Launch the User-ID Agent application. | 1. Select **User-ID Agent** from the Windows Start menu.<br> |
| Step 2 | Define the servers the User-ID agent should monitor to collect IP address to user mapping information.<br>The User-ID agent can monitor a total of up to 100 servers, including syslog senders.<br>Keep in mind that to collect all of the required mappings, you must connect to all servers that your users log in to in order to monitor the security log files on all servers that contain login events. | 1. Select **User Identification > Discovery**.<br>2. In the **Servers** section of the screen, click **Add**.<br>3. Enter a **Name** and **Server Address** for the server to be monitored. The network address can be a FQDN or an IP address.<br>4. Select the **Server Type** (**Microsoft Active Directory**, **Microsoft Exchange**, **Novell eDirectory**, or **Syslog Sender**) and then click **OK** to save the server entry. Repeat this step for each server to be monitored.<br>5. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Auto Discover**.<br>The auto-discovery locates domain controllers in the local domain only; you must manually add Exchange servers, eDirectory servers, and syslog senders.<br>6. (Optional) To tune the frequency at which the firewall polls configured servers for mapping information, select **User Identification > Setup** and **Edit** the Setup section. On the **Server Monitor** tab, modify the value in the **Server Log Monitor Frequency (seconds)** field. As a best practice, you should increase the value in this field to 5 seconds in environments with older Domain Controllers or high-latency links. Click **OK** to save the changes. |

| Map IP Addresses to Users Using the User-ID Agent (Continued) | | |
|---|---|---|
| Step 3 | (Optional) If you configured the agent to connect to a Novell eDirectory server, you must specify how the agent should search the directory. | 1. Select **User Identification > Setup** and click **Edit** in the Setup section of the window.<br>2. Select the **eDirectory** tab and then complete the following fields:<br>• **Search Base**—The starting point or root context for agent queries, for example: `dc=domain1, dc=example, dc=com`.<br>• **Bind Distinguished Name**—The account to use to bind to the directory, for example: `cn=admin, ou=IT, dc=domain1, dc=example, dc=com`.<br>• **Bind Password**—The bind account password. The agent saves the encrypted password in the configuration file.<br>• **Search Filter**—The search query for user entries (default is `objectClass=Person`).<br>• **Server Domain Prefix**—A prefix to uniquely identify the user. This is only required if there are overlapping name spaces, such as different users with the same name from two different directories.<br>• **Use SSL**—Select the check box to use SSL for eDirectory binding.<br>• **Verify Server Certificate**—Select the check box to verify the eDirectory server certificate when using SSL. |
| Step 4 | (Optional) Enable client probing.<br>Client probing is useful in environments where IP addresses are not tightly bound to users because it ensures that previously mapped addresses are still valid. However, as the total number of learned IP addresses grows, so does the amount of traffic generated. As a best practice, only enable probing on network segments where IP address turnover is high.<br>For more details on the placement of User-ID agents using client probing, refer to Architecting User Identification (User-ID) Deployments. | 1. On the **Client Probing** tab, select the **Enable WMI Probing** check box and/or the **Enable NetBIOS Probing** check box.<br>2. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client.<br> For NetBIOS probing to work effectively, each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled. WMI probing is always preferred over NetBIOS whenever possible. |
| Step 5 | Save the configuration. | Click **OK** to save the User-ID agent setup settings and then click **Commit** to restart the User-ID agent and load the new settings. |

 © Palo Alto Networks, Inc.

| Map IP Addresses to Users Using the User-ID Agent (Continued) | | |
|---|---|---|
| Step 6 | (Optional) Define the set of users for which you do not need to provide IP address-to-user name mappings, such as kiosk accounts.<br><br> You can also use the `ignore-user` list to identify users whom you want to force to authenticate using Captive Portal. | Create an `ignore_user_list.txt` file and save it to the User-ID Agent folder on the domain server where the agent is installed.<br><br>List the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Each user account name must be on a separate line. For example:<br>`SPAdmin`<br>`SPInstall`<br>`TFSReport`<br><br>You can use an asterisk as a wildcard character to match multiple usernames, but only as the last character in the entry. For example, **`corpdomain\it-admin*`** would match all administrators in the **`corpdomain`** domain whose usernames start with the string **`it-admin`**. |
| Step 7 | Configure the firewalls to connect to the User-ID agent. | Complete the following steps on each firewall you want to connect to the User-ID agent to receive user mappings:<br>1. Select **Device > User Identification > User-ID Agents** and click **Add**.<br>2. Enter a **Name** for the User-ID agent.<br>3. Enter the IP address of the Windows **Host** on which the User-ID Agent is installed.<br>4. Enter the **Port** number (1-65535) on which the agent will listen for user mapping requests. This value must match the value configured on the User-ID agent. By default, the port is set to 5007 on the firewall and on newer versions of the User-ID agent. However, some older User-ID agent versions use port 2010 as the default.<br>5. Make sure that the configuration is **Enabled**, then click **OK**.<br>6. **Commit** the changes.<br>7. Verify that the **Connected status** displays as connected (a green light). |
| Step 8 | Verify that the User-ID agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent. | 1. Launch the User-ID agent and select **User Identification**.<br>2. Verify that the agent status shows **Agent is running**. If the Agent is not running, click **Start**.<br>3. To verify that the User-ID agent can connect to monitored servers, make sure the Status for each Server is **Connected**.<br>4. To verify that the firewalls can connect to the User-ID agent, make sure the Status for each of the Connected Devices is **Connected**.<br>5. To verify that the User-ID agent is mapping IP addresses to usernames, select **Monitoring** and make sure that the mapping table is populated. You can also **Search** for specific users, or **Delete** user mappings from the list. |

## Configure User Mapping Using the PAN-OS Integrated User-ID Agent

The following procedure shows how to configure the PAN-OS integrated agent on the firewall for user mapping. The integrated User-ID agent performs the same tasks as the Windows-based agent with the exception of NetBIOS client probing (WMI probing is supported).

| Map IP Addresses to Users Using the Integrated User-ID Agent | | |
|---|---|---|
| Step 1 | Create an Active Directory (AD) account for the User-ID agent.<br><br>The account must have the privilege levels required to log in to each service or host that the User-ID agent will monitor to collect user mapping data. | • **Windows 2008 or later domain servers**—Add the account to the Event Log Readers group. If you are using the on-device User-ID agent, the account must also be a member of the Distributed COM Users Group.<br>• **Windows 2003 domain servers**—Assign Manage Auditing and Security Logs permissions through group policy.<br>• **WMI probing**—Make sure the account has rights to read the CIMV2 namespace; by default, Domain Administrator and Server Operator accounts have this permission.<br>• **NTLM authentication**—Because the firewall must join the domain if you are using Captive Portal NTLM authentication with an on-device User-ID agent, the Windows account you create for NTLM access must have administrative privileges. Note that due to AD restrictions on virtual systems running on the same host, if the firewall has multiple virtual systems, only vsys1 will be able to join the domain. |

 © Palo Alto Networks, Inc.

**Map IP Addresses to Users Using the Integrated User-ID Agent (Continued)**

| | | |
|---|---|---|
| Step 2 | Define the servers the firewall should monitor to collect IP address-to-user mapping information. You can define a total of up to 100 servers of all types, including syslog senders.<br><br>Note that to collect all the required mappings, the firewall must connect to all servers that your users log in to so it can monitor the security log files on all servers that contain login events. | 1. Select **Device > User Identification > User Mapping**.<br>2. In the Server Monitoring section of the screen, click **Add**.<br>3. Enter a **Name** and **Network Address** for the server. The network address can be a FQDN or an IP address.<br>4. Select the **Type** of server.<br>5. Make sure the **Enabled** check box is selected and then click **OK**<br>6. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Discover**.<br> <br>The auto-discovery feature is for domain controllers only; you must manually add any Exchange servers or eDirectory servers you want to monitor.<br>7. (Optional) Specify the frequency at which the firewall polls Windows servers for mapping information. This is the interval between the end of the last query and the start of the next query.<br>If the query load is high, the observed delay between queries might significantly exceed the specified frequency.<br>a. In the **Palo Alto Networks User ID Agent Setup** section, click the Edit icon.<br>b. Select the **Server Monitor** tab and specify the **Server Log Monitor Frequency** in seconds (default is 2, range is 1-3600).<br>As a best practice, increase the value in this field to 5 seconds in environments with older domain controllers or high-latency links.<br>c. Click **OK** to save the changes. |
| Step 3 | Set the domain credentials for the account the firewall will use to access Windows resources. This is required for monitoring Exchange servers and domain controllers as well as for WMI probing. | 1. Edit the Palo Alto Networks User ID Agent Setup section of the screen.<br>2. On the **WMI Authentication** tab, enter the **User Name** and **Password** for the account that will be used to probe the clients and monitor servers. Enter the user name using the domain\username syntax.<br> |

| Map IP Addresses to Users Using the Integrated User-ID Agent (Continued) | | |
|---|---|---|
| Step 4 | (Optional) Enable WMI probing.<br> The on-device agent does not support NetBIOS probing; it is supported on the Windows-based User-ID agent only. | 1. On the **Client Probing** tab, select the **Enable Probing** check box.<br>2. (Optional) If necessary, modify the **Probe Interval** (in minutes) to ensure it is long enough for the User-ID agent to probe all the learned IP addresses (default is 20, range is 1-1440). This is the interval between the end of the last probe request and the start of the next request.<br> If the request load is high, the observed delay between requests might significantly exceed the specified interval.<br>3. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client. |
| Step 5 | Save the configuration. | 1. Click **OK** to save the User-ID agent setup settings.<br>2. Click **Commit** to save the configuration. |
| Step 6 | (Optional) Define the set of users for which you do not need to provide IP address-to-user name mappings, such as kiosk accounts.<br> You can also use the `ignore-user` list to identify users whom you want to force to authenticate using Captive Portal. | 1. Open a CLI session to the firewall.<br>2. To add the list of user accounts for which you do not want the firewall to perform mapping, run the following command:<br>**set user-id-collector ignore-user <value>**<br>where <value> is a list of the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Separate entries with a space and do not include the domain name with the username. For example:<br>`set user-id-collector ignore-user SPAdmin SPInstall TFSReport`<br>3. **Commit** your changes. |
| Step 7 | Verify the configuration. | 1. From the CLI, enter the following command:<br>`show user server-monitor state all`<br>2. On the **Device > User Identification > User Mapping** tab in the web interface, verify that the **Status** of each server you configured for server monitoring is **Connected**. |


© Palo Alto Networks, Inc.

# Configure User-ID to Receive User Mappings from a Syslog Sender

The following topics describe how to configure the User-ID agent (either the Windows agent or the integrated agent on the firewall) as a Syslog listener:

- Configure the Integrated User-ID Agent as a Syslog Listener
- Configure the Windows User-ID Agent as a Syslog Listener

## Configure the Integrated User-ID Agent as a Syslog Listener

The following workflow describes how to configure the PAN-OS integrated User-ID agent to receive syslog messages from authenticating services.

The PAN-OS integrated User-ID agent accepts syslogs over SSL and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages. However, if you must use UDP, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending UDP traffic to the firewall.

| **Collect User Mappings from Syslog Senders** | |
|---|---|
| Step 1 Determine whether there is a pre-defined syslog filter for your particular syslog sender(s).<br><br>Palo Alto Networks provides several pre-defined syslog filters, which are delivered as Application content updates and are therefore updated dynamically as new filters are developed. The pre-defined filters are global to the firewall, whereas manually-defined filters apply to a single virtual system only.<br><br> Any new syslog filters in a given content release will be documented in the corresponding release note along with the specific regex used to define the filter. | 1. Verify that your Application or Application and Threat database is up to date:<br>a. Select **Device > Dynamic Updates**.<br>b. Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates.<br>c. If a new update is available, **Download** and **Install** it.<br>2. Check to see what pre-defined filters are available:<br>a. Select **Device > User Identification > User Mapping**.<br>b. In the Server Monitoring section of the screen, click **Add**.<br>c. Select **Syslog Sender** as the server **Type**.<br>d. Select the **Filter** drop-down and check to see if there is a filter for the manufacturer and product you plan to forward syslogs from. If the filter you need is available, skip to Step 5 for instructions on defining the servers. If the filter you need is not available, continue to Step 2. |

**Collect User Mappings from Syslog Senders (Continued)**

Step 2 Manually define syslog filter(s) for extracting the User-ID IP address to username mapping information from syslog messages.

In order to be parsed by the User-ID agent, syslog messages must meet the following criteria:

- Each syslog message must be a single-line text string. Line breaks are delimited by a carriage return and a new line (`\r\n`) or a new line (`\n`).
- The maximum allowed size of an individual syslog message is 2048 bytes.
- Syslog messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets.
- A single packet may contain multiple syslog messages.

1. Review the syslogs generated by the authenticating service to identify the syntax of the login events. This enables you to create the matching patterns that will allow the firewall to identify and extract the authentication events from the syslogs.

   While reviewing the syslogs, also determine whether the domain name is included in the log entries. If the authentication logs do not contain domain information, consider defining a default domain name when adding the syslog sender to the monitored servers list in Step 5.
2. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User-ID Agent Setup section.
3. On the **Syslog Filters** tab, **Add** a new syslog parse profile.
4. Enter a name for the **Syslog Parse Profile**.
5. Specify the **Type** of parsing to use to filter out the user mapping information by selecting one of the following options:
   - **Regex Identifier**—With this type of parsing, you specify regular expressions to describe search patterns for identifying and extracting user mapping information from syslog messages. Continue to Step 3 for instructions on creating the regex identifiers.
   - **Field Identifier**—With this type of parsing, you specify a string to match the authentication event, and prefix and suffix strings to identify the user mapping information in the syslogs. Continue to Step 4 for instructions on creating the field identifiers.

 © Palo Alto Networks, Inc.

**Collect User Mappings from Syslog Senders (Continued)**

Step 3 If you selected **Regex Identifier** as the parsing **Type**, create the regex matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a regex configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter, you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event Regex** field. For example, when matched against the example syslog message, the following regex instructs the firewall to extract the first `{1}` instance of the string `authentication success`. The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character: `(authentication\ success){1}`.
2. Enter the regex for identifying the beginning of the username in the authentication success messages in the **Username Regex** field. For example, the regex `User:([a-zA-Z0-9\\\._]+)` would match the string `User:johndoe1` in the example message and extract `acme\johndoe1` as the User-ID.

   If the syslogs do not contain domain information and you require domain names in your user mappings, be sure to enter the **Default Domain Name** when defining the monitored server entry in Step 5.
3. Enter the regex for identifying the IP address portion of the authentication success messages in the **Address Regex** field. For example, the following regular expression `Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})` would match an IPv4 address (`Source:192.168.0.212` in the example syslog).
4. Click **OK**.

**Collect User Mappings from Syslog Senders (Continued)**

Step 4 If you selected **Field Identifier** as the parsing **Type**, define the string matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a field identifier configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter, you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event String** field. For example, when matched against the sample syslog message, you would enter the string `authentication success` to identify authentication events in the syslog.
2. Enter the matching string for identifying the beginning of the username field within the authentication syslog message in the **Username Prefix** field. For example, the string `User:` identifies the beginning of the username field in the sample syslog.
3. Enter the **Username Delimiter** to mark the end of the username field within an authentication syslog message. For example, if the username is followed by a space, you would enter `\s` to indicate that the username field is delimited by a standalone space in the sample log.
4. Enter the matching string for identifying the beginning of the IP address field within the authentication event log in the **Address Prefix** field. For example, the string `Source:` identifies the beginning of the address field in the example log.
5. Enter the **Address Delimiter** to mark the end of the IP address field within the authentication success message within the field. For example, if the address is followed by a line break, you would enter `\n` to indicate that the address field is delimited by a new line.
6. Click **OK**.

Step 5 Define the servers that will send syslog messages to the firewall for user mapping purposes.

You can define a total of up to 100 monitored servers of all types, including syslog senders. The firewall will discard any syslog messages received from servers that are not on this list.

A Syslog sender using SSL to connect will only show a **Status** of **Connected** when there is an active SSL connection. Syslog senders using UDP will not show a **Status** value.

1. Select **Device > User Identification > User Mapping**.
2. In the Server Monitoring section of the screen, click **Add**.
3. Enter a **Name** and **Network Address** for the server.
4. Select **Syslog Sender** as the server **Type**.
5. Make sure the **Enabled** check box is selected.
6. (Optional) If the syslogs that the authenticating device sends do not include domain information in the login event logs, enter the **Default Domain Name** to append to the user mappings.
7. Click **OK** to save the settings.

 © Palo Alto Networks, Inc.

| Collect User Mappings from Syslog Senders (Continued) | | |
|---|---|---|
| Step 6 | Enable syslog listener services in the management profile associated with the interface used for user mapping.<br> <br>Use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages when using agentless User Mapping on a firewall. However, if you must use UDP, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending UDP traffic to the firewall. | 1. Select **Network > Network Profiles > Interface Mgmt** and then select an Interface Management profile to edit or click **Add** to create a new profile.<br>2. Select **User-ID Syslog Listener-SSL** and/or **User-ID Syslog Listener-UDP**, depending on the protocols you defined when you set up your syslog senders in the Server Monitor list.<br> <br>On the Windows User-ID agent, the default listening port for syslog over UDP or TCP is 514, but the port value is configurable. For the agentless User Mapping feature on the firewall, only syslog over UDP and SSL are supported and the listening ports (514 for UDP and 6514 for SSL) are not configurable; they are enabled through the management service only.<br>3. Click **OK** to save the interface management profile.<br> <br>Even after enabling the User-ID Syslog Listener service on the interface, the interface will only accept syslog connections from servers that have a corresponding entry in the User-ID monitored servers configuration. Connections or messages from servers that are not on the list will be discarded. |
| Step 7 | Save the configuration. | Click **Commit** to save the configuration. |

**Collect User Mappings from Syslog Senders (Continued)**

Step 8 Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
        UDP Syslog Listener Service is enabled
        SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)     Host: Syslog2(10.5.204.41)
        number of log messages                          : 1000
        number of auth. success messages                : 1000
        number of active connections                    : 0
        total connections made                          : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name                            TYPE     Host            Vsys    Status

------------------------------------------------------------------------------
AD                              AD       10.2.204.43     vsys1   Connected

Syslog Servers:
Name                      Connection Host            Vsys    Status
------------------------------------------------------------------------------
Syslog1                   UDP        10.5.204.40     vsys1   N/A
Syslog2                   SSL        10.5.204.41     vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP              Vsys   From    User                              IdleTimeout(s) M
axTimeout(s)
--------------- ------ ------- --------------------------------- -------------- -
192.168.3.8     vsys1  SYSLOG  acme\jreddick                     2476           2
476
192.168.5.39    vsys1  SYSLOG  acme\jdonaldson                   2480           2
480
192.168.2.147   vsys1  SYSLOG  acme\ccrisp                       2476           2
476
192.168.2.175   vsys1  SYSLOG  acme\jjaso                        2476           2
476
192.168.4.196   vsys1  SYSLOG  acme\jblevins                     2480           2
480
192.168.4.103   vsys1  SYSLOG  acme\bmoss                        2480           2
480
192.168.2.193   vsys1  SYSLOG  acme\esogard                      2476           2
476
192.168.2.119   vsys1  SYSLOG  acme\acallaspo                    2476           2
476
192.168.3.176   vsys1  SYSLOG  acme\jlowrie                      2478           2
478

Total: 9 users
```

## Configure the Windows User-ID Agent as a Syslog Listener

The following workflow describes how to configure a Windows-based User-ID agent to listen for syslogs from authenticating services.

The Windows User-ID agent accepts syslogs over TCP and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, use TCP instead of UDP. In either case, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

 © Palo Alto Networks, Inc.

| Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders | |
|---|---|
| Step 1 Launch the User-ID Agent application. | 1. Click Start and select **User-ID Agent**.<br> |
| Step 2 Manually define syslog filter(s) for extracting the User-ID IP address to username mapping information from syslog messages.<br><br>In order to be parsed by the User-ID agent, syslog messages must meet the following criteria:<br>• Each syslog message must be a single-line text string. Line breaks are delimited by a carriage return and a new line (\r\n) or a new line (\n).<br>• The maximum allowed size of an individual syslog message is 2048 bytes.<br>• Syslog messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets.<br>• A single packet may contain multiple syslog messages. | 1. Review the syslogs generated by the authenticating service to identify the syntax of the login events. This enables you to create the matching patterns that will allow the firewall to identify and extract the authentication events from the syslogs.<br> While reviewing the syslogs, also determine whether the domain name is included in the log entries. If the authentication logs do not contain domain information, consider defining a default domain name when adding the syslog sender to the monitored servers list in Step 5.<br>2. Select **User Identification > Setup** and click **Edit** in the Setup section of the dialog.<br>3. On the **Syslog** tab, **Add** a new syslog parse profile.<br>4. Enter a **Profile Name** and **Description**.<br>5. Specify the **Type** of parsing to use to filter out the user mapping information by selecting one of the following options:<br>• **Regex**—With this type of parsing, you specify regular expressions to describe search patterns for identifying and extracting user mapping information from syslog messages. Continue to Step 3 for instructions on creating the regex identifiers.<br>• **Field**—With this type of parsing, you specify a sting to match the authentication event, and prefix and suffix strings to identify the user mapping information in the syslogs. Continue to Step 4 for instructions on creating the field identifiers. |

**Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)**

Step 3 If you selected **Regex** as the parsing **Type**, create the regex matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a regex configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter, you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event Regex** field. For example, when matched against the example syslog message, the following regex instructs the firewall to extract the first `{1}` instance of the string `authentication success`. The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character: `(authentication\ success){1}`.
2. Enter the regex for identifying the beginning of the username in the authentication success messages in the **Username Regex** field. For example, the regex `User:([a-zA-Z0-9\\\._]+)` would match the string `User:johndoe1` in the example message and extract `acme\johndoe1` as the User-ID.

   If the syslogs do not contain domain information and you require domain names in your user mappings, be sure to enter the **Default Domain Name** when defining the monitored server entry in Step 5.
3. Enter the regex for identifying the IP address portion of the authentication success messages in the **Address Regex** field. For example, the following regular expression `Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})` would match an IPv4 address (`Source:192.168.0.212` in the example syslog).
4. Click **OK** to save the profile.

 © Palo Alto Networks, Inc.

**Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)**

**Step 4** If you selected **Field Identifier** as the parsing **Type**, define the string matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a field identifier configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```

If the syslog contains a standalone space and/or tab as a delimiter, you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event String** field. For example, when matched against the sample syslog message, you would enter the string `authentication success` to identify authentication events in the syslog.
2. Enter the matching string for identifying the beginning of the username field within the authentication syslog message in the **Username Prefix** field. For example, the string `User:` identifies the beginning of the username field in the sample syslog.
3. Enter the **Username Delimiter** to mark the end of the username field within an authentication syslog message. For example, if the username is followed by a space, you would enter `\s` to indicate that the username field is delimited by a standalone space in the sample log.
4. Enter the matching string for identifying the beginning of the IP address field within the authentication event log in the **Address Prefix** field. For example, the string `Source:` identifies the beginning of the address field in the example log.
5. Enter the **Address Delimiter** to mark the end of the IP address field within the authentication success message within the field. For example, if the address is followed by a line break, you would enter `\n` to indicate that the address field is delimited by a new line.
6. Click **OK** to save the profile.

**Step 5** Enable the syslog listening service on the agent.

As a best practice, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

1. Select the **Enable Syslog Service** check box.
2. (Optional) Modify the **Syslog Service Port** number to match the port number used by the syslog sender (default=514).
3. To save the agent syslog configuration, click **OK**.

**Step 6** Define the servers that will send syslog messages to the User-ID agent.

You can define up to 100 syslog senders. The User-ID agent will discard any syslog messages received from servers that are not on this list.

1. Select **User Identification > Discovery**.
2. In the **Servers** section of the screen, click **Add**.
3. Enter a **Name** and **Server Address** for the server that will send syslogs to the agent.
4. Select **Syslog Sender** as the **Server Type**.
5. Select a **Filter** you defined in Step 2.
6. (Optional) If the syslogs that the authenticating device sends do not include domain information in the login event logs, enter the **Default Domain Name** to append to the user mappings.
7. Click **OK** to save the settings.

**Step 7** Save the configuration.

Click **Commit** to save the configuration.

**Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)**

Step 8 Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
        UDP Syslog Listener Service is enabled
        SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)     Host: Syslog2(10.5.204.41)
        number of log messages                          : 1000
        number of auth. success messages                : 1000
        number of active connections                    : 0
        total connections made                          : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name                            TYPE     Host            Vsys    Status

------------------------------------------------------------------------------
AD                              AD       10.2.204.43     vsys1   Connected

Syslog Servers:
Name                      Connection Host            Vsys    Status
------------------------------------------------------------------------------
Syslog1                   UDP        10.5.204.40     vsys1   N/A
Syslog2                   SSL        10.5.204.41     vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP              Vsys   From    User                              IdleTimeout(s) M
axTimeout(s)
--------------- ------ ------- --------------------------------- -------------- -
192.168.3.8     vsys1  SYSLOG  acme\jreddick                     2476           2
476
192.168.5.39    vsys1  SYSLOG  acme\jdonaldson                   2480           2
480
192.168.2.147   vsys1  SYSLOG  acme\ccrisp                       2476           2
476
192.168.2.175   vsys1  SYSLOG  acme\jjaso                        2476           2
476
192.168.4.196   vsys1  SYSLOG  acme\jblevins                     2480           2
480
192.168.4.103   vsys1  SYSLOG  acme\bmoss                        2480           2
480
192.168.2.193   vsys1  SYSLOG  acme\esogard                      2476           2
476
192.168.2.119   vsys1  SYSLOG  acme\acallaspo                    2476           2
476
192.168.3.176   vsys1  SYSLOG  acme\jlowrie                      2478           2
478

Total: 9 users
```

 © Palo Alto Networks, Inc.

# Map IP Addresses to Usernames Using Captive Portal

If the firewall receives a request from a security zone that has User-ID enabled and the source IP address does not have any user data associated with it yet, the firewall checks its Captive Portal policy rules for a match to determine whether to perform authentication. This is useful in environments where you have clients that are not logged in to your domain servers, such as Linux clients. The firewall triggers this user mapping method only for web traffic (HTTP or HTTPS) that matches a Captive Portal rule but has not been mapped using a different method.

- Captive Portal Authentication Methods
- Captive Portal Modes
- Configure Captive Portal

## Captive Portal Authentication Methods

Captive Portal uses the following methods to obtain user information from the client when a web request matches a Captive Portal rule:

| Authentication Method | Description |
|---|---|
| Kerberos SSO | The firewall uses Kerberos Single Sign-On (SSO) to transparently obtain user credentials. To use this method, your network requires a Kerberos infrastructure, including a key distribution center (KDC) with an authentication server and ticket granting service. The firewall must have a Kerberos account, including a principal name and password.<br>If Kerberos SSO authentication fails, the firewall falls back to NT LAN Manager (NTLM) authentication. If you don't configure NTLM, or NTLM authentication fails, the firewall falls back to web form or client certificate authentication, depending on your Captive Portal configuration. |
| NT LAN Manager (NTLM) | The firewall uses an encrypted challenge-response mechanism to obtain the user credentials from the browser. When configured properly, the browser will transparently provide the credentials to the firewall without prompting the user, but will prompt for credentials if necessary.<br>If you use the Windows-based User-ID agent, NTLM responses go directly to the domain controller where you installed the agent.<br>If you configure Kerberos SSO authentication, the firewall tries that method first before falling back to NTLM authentication. If the browser can't perform NTLM or if NTLM authentication fails, the firewall falls back to web form or client certificate authentication, depending on your Captive Portal configuration.<br>Microsoft Internet Explorer supports NTLM by default. You can configure Mozilla Firefox and Google Chrome to also use NTLM but you can't use NTLM to authenticate non-Windows clients. |

| Authentication Method | Description |
|---|---|
| Web Form | The firewall redirects web requests to a web form for authentication. You can configure Captive Portal to use a local user database, RADIUS server, TACACS+ server, LDAP server, or Kerberos server to authenticate users. Although the firewall always prompts users for credentials, this method works with all browsers and operating systems. |
| Client Certificate Authentication | The firewall prompts the browser to present a valid client certificate to authenticate the user. To use this method, you must provision client certificates on each user system and install the trusted certificate authority (CA) certificate used to issue those certificates on the firewall. |

## Captive Portal Modes

The Captive Portal mode defines how the firewall captures web requests for authentication:

| Mode | Description |
|---|---|
| Transparent | The firewall intercepts the browser traffic per the Captive Portal rule and impersonates the original destination URL, issuing an HTTP 401 to invoke authentication. However, because the firewall does not have the real certificate for the destination URL, the browser displays a certificate error to users attempting to access a secure site. Therefore, you should only use this mode when absolutely necessary, such as in Layer 2 or virtual wire deployments. |
| Redirect | The firewall intercepts unknown HTTP or HTTPS sessions and redirects them to a Layer 3 interface on the firewall using an HTTP 302 redirect to perform authentication. This is the preferred mode because it provides a better end-user experience (no certificate errors). However, it does require additional Layer 3 configuration. Another benefit of the Redirect mode is that it provides for the use of session cookies, which enable the user to continue browsing to authenticated sites without requiring re-mapping each time the time outs expire. This is especially useful for users who roam from one IP address to another (for example, from the corporate LAN to the wireless network) because they won't need to re-authenticate when the IP address changes as long as the session stays open.<br>If you use Kerberos SSO or NTLM authentication, you must use Redirect mode because the browser will provide credentials only to trusted sites. |

## Configure Captive Portal

The following procedure shows how to configure Captive Portal using the PAN-OS integrated User-ID agent to redirect web requests that match a Captive Portal rule to a redirect host. A redirect host is the intranet hostname (a hostname with no period in its name) that resolves to the IP address of the Layer 3 interface on the firewall to which the firewall will redirect requests.

If you use Captive Portal without the other User-ID functions (user mapping and group mapping), you don't need to configure a User-ID agent.

 © Palo Alto Networks, Inc.

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent | | |
|---|---|---|
| Step 1 | Configure the interfaces that the firewall will use for redirecting web requests, authenticating users, and communicating with directory servers to map usernames to IP addresses.<br><br>The firewall uses the management (MGT) interface for all these functions by default, but you can configure other interfaces. In redirect mode, you must use a Layer 3 interface for redirecting requests. | 1. (MGT interface only) Select **Device > Setup > Management**, edit the Management Interface Settings, select the **User ID** check box, and click **OK**.<br>2. (Non-MGT interface only) Assign an Interface Management profile to the Layer 3 interface that the firewall will use to redirect web requests and/or communicate with directory servers.<br>a. Select **Network > Network Profiles > Interface Mgmt** and **Add** an Interface Management profile with a unique **Name**.<br>b. Select the **Response Pages** and **User ID** check boxes, and click **OK**.<br>c. Select **Network > Interfaces**, click the interface, select the **Advanced** tab, select the Interface **Management Profile** you just created, and click **OK**.<br>3. (Non-MGT interface only) Configure a service route for the interface that the firewall will use to authenticate users. If the firewall has more than one virtual system (vsys), the service route can be global or vsys-specific. The services must include **LDAP** and potentially the following:<br>– **Kerberos**, **RADIUS**, or **TACACS+**—Configure a service route for one of these services only if you will use it for external authentication.<br>– **UID Agent**—Configure this service only if you will enable NT LAN Manager (NTLM) authentication or if you will Configure a Firewall to Share User Mapping Data with Other Firewalls.<br>4. (Redirect mode only) Create a DNS address (A) record that maps the IP address on the Layer 3 interface to the redirect host. If you will use Kerberos SSO, you must also add a DNS pointer (PTR) record that performs the same mapping.<br><br>If your network doesn't support access to the directory servers from any firewall interface, you must Configure User Mapping Using the Windows User-ID Agent. |
| Step 2 | Make sure Domain Name System (DNS) is configured to resolve your domain controller addresses. | To verify proper resolution, ping the server FQDN. For example:<br>`admin@PA-200> ping host dc1.acme.com` |
| Step 3 | Create a Kerberos keytab for the redirect host.<br><br>Required for Kerberos SSO authentication. | Create a Kerberos keytab. A keytab is a file that contains Kerberos account information (principal name and hashed password) for the redirect host (the firewall).<br><br>To support Kerberos SSO, your network must have a Kerberos infrastructure, including a key distribution center (KDC) with an authentication server and ticket granting service. |

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | |
|---|---|
| Step 4 Configure clients to trust Captive Portal certificates.<br><br>Required for redirect mode—to transparently redirect users without displaying certificate errors. You can generate a self-signed certificate or import a certificate that an external certificate authority (CA) signed. | To use a self-signed certificate, create a root CA certificate and use it to sign the certificate you will use for Captive Portal:<br>1. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Create a Self-Signed Root CA Certificate or import a CA certificate (see Import a Certificate and Private Key).<br>3. Generate a Certificate on the Device to use for Captive Portal. Be sure to configure the following fields:<br>• **Common Name**—Enter the DNS name of the intranet host for the Layer 3 interface.<br>• **Signed By**—Select the CA certificate you just created or imported.<br>• Certificate Attributes—Click **Add**, for the **Type** select **IP** and, for the **Value**, enter the IP address of the Layer 3 interface to which the firewall will redirect requests.<br>4. Configure an SSL/TLS Service Profile. Assign the Captive Portal certificate you just created to the profile.<br>5. Configure clients to trust the certificate:<br>a. Export the CA certificate you created or imported.<br>b. Import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory (AD) Group Policy Object (GPO). |

 © Palo Alto Networks, Inc.

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | | |
|---|---|---|
| Step 5 | Configure an authentication server profile.<br><br>Required for external authentication. If you enable Kerberos SSO or NTLM authentication, the firewall uses the external service only if those methods fail. | • Configure a RADIUS Server Profile.<br>• Configure a TACACS+ Server Profile<br>• Configure an LDAP Server Profile<br>• Configure a Kerberos Server Profile<br><br>The PAN-OS web server timeout (default is 3 seconds) must be the same as or greater than the server profile timeout multiplied by the number of servers in the profile. For RADIUS and TACACS+, the default server profile **Timeout** is 3 seconds. For LDAP, the timeout is the total of the **Bind Timeout** (default is 30 seconds) and **Search Timeout** (default is 30 seconds) for each server. For Kerberos, the non-configurable timeout can take up to 17 seconds for each server. Also, the Captive Portal session timeout (default is 30 seconds) must be greater than the web server timeout.<br><br>To change the web server timeout, enter the following firewall CLI command, where **<value>** is 3-30 seconds: **`set deviceconfig setting l3-service timeout <value>`**. To change the Captive Portal session timeout, select **Device > Setup > Session**, edit the Session Timeouts, and enter a new **Captive Portal** value in seconds (range is 1-1,599,999). Keep in mind that the more you raise the web server and Captive Portal session timeouts, the slower Captive Portal will respond to users. |
| Step 6 | Add the users and user groups to the local database on the firewall.<br><br>Required for local database authentication. If you enable Kerberos SSO and/or NTLM authentication, the firewall uses the local database only if those methods fail. | Create the local database. Add the users who will authenticate using Captive Portal. |

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | |
|---|---|
| Step 7 Add an authentication profile<br><br>The profile defines the authentication methods to use (Kerberos SSO, external service, or local database) when a Captive Portal rule invokes Web Form authentication. Even if you enable NTLM, you must define a secondary authentication method in case NTLM authentication fails or the User-ID agent doesn't support NTLM.<br><br> If you set the authentication **Type** to **RADIUS**, specify a RADIUS **User Domain** in case users don't enter the domain at login. | Configure an authentication profile:<br>1. If the authentication **Type** is an external service (**RADIUS**, **TACACS+**, **LDAP**, or **Kerberos**), select the authentication **Server Profile** you created.<br>2. If you use Kerberos SSO, enter the **Kerberos Realm** (usually the DNS domain of the users, except that the realm is uppercase), and import the **Kerberos Keytab** you created.<br>3. In the Allow List on the **Advanced** tab, **Add** the users and user groups that can authenticate using this profile. If the authentication **Type** is **Local Database**, add the Captive Portal users or user groups you created. You can select **all** to allow every user to authenticate. After you complete the Allow List, click **OK**.<br><br> If your users are in multiple domains or Kerberos realms, you can create an authentication profile for each domain or realm, assign all the profiles to the authentication sequence (Step 8), and assign the sequence to the Captive Portal configuration. |
| Step 8 (Optional) Add an authentication sequence<br><br> If the firewall might try multiple authentication profiles in the sequence for any one user (for example, if some directory server connections are unreliable), then the PAN-OS web server timeout must be the same as or greater than the timeout for the sequence, which is the total of the timeouts for all its authentication profiles. Also, the session timeout for Captive Portal must be greater than the web server timeout. To change these timeouts, see the note in Step 5. | Configure an authentication sequence:<br>1. Select **Device > Authentication Sequence**, click **Add**, and enter a **Name** to identify the authentication sequence.<br>2. Select the **Use domain to determine authentication profile** check box: the device will match the domain name that a user enters during login with the **User Domain** or **Kerberos Realm** of an authentication profile in the sequence, and then use that profile to authenticate the user.<br>3. For each authentication profile to include, click **Add** and select the profile from the drop-down.<br>4. Click **OK** to save the authentication sequence. |

 © Palo Alto Networks, Inc.

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | |
|---|---|
| Step 9 Configure Client Certificate Authentication.<br><br>Required if Captive Portal will use this authentication method.<br><br>You don't need an authentication profile or sequence for client certificate authentication. If you configure both an authentication profile/sequence and certificate authentication, users must authenticate using both. | 1. Use a root CA certificate to generate a client certificate for each user who will authenticate to Captive Portal. The CA in this case is usually your enterprise CA, not the firewall.<br>2. Export the CA certificate in PEM format to a system that the firewall can access.<br>3. Import the CA certificate onto the firewall: see Import a Certificate and Private Key. After the import, click the imported certificate, select **Trusted Root CA**, and click **OK**.<br>4. Configure a Certificate Profile.<br>• In the **Username Field** drop-down, select the certificate field that contains the user identity information.<br>• In the **CA Certificates** list, click **Add** and select the CA certificate you just imported. |
| Step 10 Enable NT LAN Manager (NTLM) authentication.<br><br>Required for NLTM authentication.<br><br>When using the PAN-OS integrated User-ID agent, the firewall must successfully resolve the DNS name of your domain controller to join the domain (using the credentials you enter in this step). | 1. If you haven't already done so, Create an Active Directory (AD) account for the User-ID agent.<br>2. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User ID Agent Setup section.<br>3. On the **NTLM** tab, select the **Enable NTLM authentication processing** check box.<br>4. Enter the **NTLM Domain** against which the User-ID agent on the firewall will check NTLM credentials.<br>5. In the **Admin User Name**, **Password**, and **Confirm Password** fields, enter the username and password of the Active Directory account you created for the User-ID agent.<br>Do not include the domain in the **Admin User Name** field. Otherwise, the firewall will fail to join the domain.<br>Palo Alto Networks recommends that you use a User-ID agent account that is separate from your firewall administrator account.<br>6. You don't need to configure any other settings for the User-ID agent: click **OK**. |

| Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued) | |
|---|---|
| Step 11 Configure the Captive Portal settings. | 1. Select **Device > User Identification > Captive Portal Settings** and edit the settings.<br>2. Make sure the **Enable Captive Portal** check box is selected.<br>3. Select the **SSL/TLS Service Profile** you created for redirect requests over TLS.<br>4. Select the **Mode** (in this example, **Redirect**).<br>5. (Redirect mode only) Specify the **Redirect Host** name that resolves to the IP address of the Layer 3 interface for redirected requests.<br>6. Select the authentication method to use if NTLM fails (or if you don't use NTLM):<br>• To use Kerberos SSO, an external server, or the local database, select the **Authentication Profile** or authentication sequence you created.<br>• To use client certificate authentication, select the **Certificate Profile** you created.<br>7. Click **OK** and **Commit** to save the Captive Portal configuration. |

 © Palo Alto Networks, Inc.

# Configure User Mapping for Terminal Server Users

Individual terminal server users appear to have the same IP address and therefore an IP address-to-username mapping is not sufficient to identify a specific user. To enable identification of specific users on Windows-based terminal servers, the Palo Alto Networks Terminal Services agent (TS agent) allocates a port range to each user. It then notifies every connected firewall about the allocated port range, which allows the firewall to create an IP address-port-user mapping table and enable user- and group-based security policy enforcement. For non-Windows terminal servers, you can configure the User-ID XML API to extract user mapping information.

The following sections describe how to configure user mapping for terminal server users:

- Configure the Palo Alto Networks Terminal Server Agent for User Mapping
- Retrieve User Mappings from a Terminal Server Using the User-ID XML API

## Configure the Palo Alto Networks Terminal Server Agent for User Mapping

Use the following procedure to install the TS agent on the terminal server. You must install the TS agent on all terminal servers that your users log in to in order to successfully map all your users.

For information about the supported terminal servers supported by the TS Agent, refer to "Operating System (OS) Compatibility TS Agent" in the Terminal Services Agent Release Notes, which are available on the Palo Alto Networks Software Updates page.

| **Install the Windows Terminal Server Agent** | |
|---|---|
| Step 1 Download the TS Agent installer. | 1. Log in to the Palo Alto Networks Support site.<br>2. Select **Software Updates** from the Manage Devices section.<br>3. Scroll to the **Terminal Services Agent** section and **Download** the version of the agent you want to install.<br>4. Save the `TaInstall64.x64-x.x.x-xx.msi` or `TaInstall-x.x.x-xx.msi` file (be sure to select the appropriate version based on whether the Windows system is running a 32-bit OS or a 64-bit OS) on the system(s) where you plan to install the agent. |

**Install the Windows Terminal Server Agent (Continued)**

| | | |
|---|---|---|
| Step 2 | Run the installer as an administrator.<br> | 1. To launch a command prompt as an administrator, click Start and right-click **Command Prompt** and then select **Run as administrator**.<br>2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop you would enter the following:<br>`C:\Users\administrator.acme>cd Desktop`<br>`C:\Users\administrator.acme\Desktop>TaInstall-6.0.0-1.msi`<br>3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to the `C:\Program Files (x86)\Palo Alto Networks\Terminal Server Agent` folder, but you can **Browse** to a different location.<br>4. When the installation completes, **Close** the setup window.<br>If you are upgrading to a TS Agent version that has a newer driver than the existing installation, the installation wizard prompts you to reboot the system after upgrading in order to use the new driver. |
| Step 3 | Launch the Terminal Server Agent application. | Click Start and select **Terminal Server Agent**.<br> 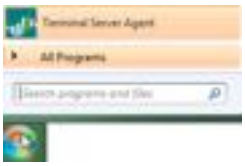 |

 © Palo Alto Networks, Inc.

**Install the Windows Terminal Server Agent (Continued)**

Step 4 Define the range of ports for the TS Agent to allocate to end users.

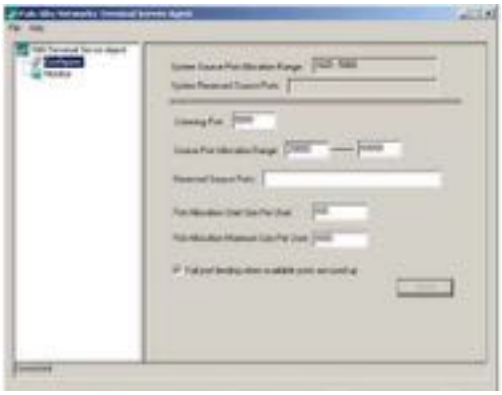

The **System Source Port Allocation Range** and **System Reserved Source Ports** fields specify the range of ports that will be allocated to non-user sessions. Make sure the values specified in these fields do not overlap with the ports you designate for user traffic. These values can only be changed by editing the corresponding Windows registry settings.

1. Select **Configure**.
2. Set the **Source Port Allocation Range** (default 20000-39999). This is the full range of port numbers that the TS Agent will allocate for user mapping. The port range you specify cannot overlap with the **System Source Port Allocation Range**.
3. (Optional) If there are ports/port ranges within the source port allocation that you do not want the TS Agent to allocate to user sessions, specify them as **Reserved Source Ports**. To include multiple ranges, use commas with no spaces, for example: `2000-3000,3500,4000-5000`.
4. Specify the number of ports to allocate to each individual user upon login to the terminal server in the **Port Allocation Start Size Per User** field (default 200).
5. Specify the **Port Allocation Maximum Size Per User**, which is the maximum number of ports the Terminal Server agent can allocate to an individual user.
6. Specify whether to continue processing traffic from the user if the user runs out of allocated ports. By default, the **Fail port binding when available ports are used up** is selected, which indicates that the application will fail to send traffic when all ports are used. To enable users to continue using applications when they run out of ports, clear this check box. Keep in mind that this traffic may not be identified with User-ID.

Step 5 Configure the firewalls to connect to the Terminal Server agent.

Complete the following steps on each firewall you want to connect to the Terminal Server agent to receive user mappings:

1. Select **Device > User Identification > Terminal Server Agents** and click **Add**.
2. Enter a **Name** for the Terminal Server agent.
3. Enter the IP address of the Windows **Host** on which the Terminal Server agent is installed.
4. Enter the **Port** number on which the agent will listen for user mapping requests. This value must match the value configured on the Terminal Server agent. By default, the port is set to 5009 on the firewall and on the agent. If you change it here, you must also change the **Listening Port** field on the Terminal Server agent **Configure** screen.
5. Make sure that the configuration is **Enabled** and then click **OK**.
6. **Commit** the changes.
7. Verify that the **Connected** status displays as connected (a green light).

| Install the Windows Terminal Server Agent (Continued) | |
|---|---|
| Step 6 Verify that the Terminal Server agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.<br> 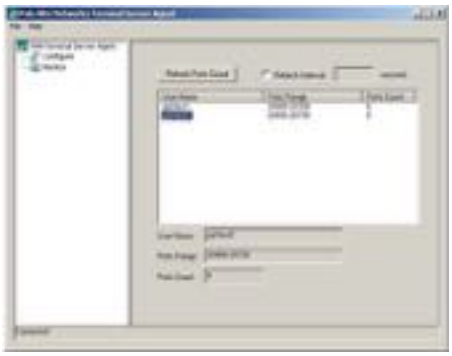 | 1. Launch the Terminal Server agent and verify that the firewalls can connect by making sure the **Connection Status** for each of Device in the Connection List is **Connected**.<br>2. To verify that the Terminal Server agent is successfully mapping port ranges to usernames, select **Monitoring** and make sure that the mapping table is populated. |

## Retrieve User Mappings from a Terminal Server Using the User-ID XML API

The User-ID XML API is a RESTful API that uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services.

To enable a non-Windows terminal server to send user mapping information directly to the firewall, create scripts that extract the user login and logout events and use them for input to the User-ID XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget and providing the firewall's API key for secure communication. Creating user mappings from multi-user systems such as terminal servers requires use of the following API messages:

- **<multiusersystem>**—Sets up the configuration for an XML API Multi-user System on the firewall. This message allows for definition of the terminal server IP address (this will be the source address for all users on that terminal server). In addition, the `<multiusersystem>` setup message specifies the range of source port numbers to allocate for user mapping and the number of ports to allocate to each individual user upon login (called the *block size*). If you want to use the default source port allocation range (1025-65534) and block size (200), you do not need to send a `<multiusersystem>` setup event to the firewall. Instead, the firewall will automatically generate the XML API Multi-user System configuration with the default settings upon receipt of the first user login event message.
- **<blockstart>**—Used with the `<login>` and `<logout>` messages to indicate the starting source port number allocated to the user. The firewall then uses the block size to determine the actual range of port numbers to map to the IP address and username in the login message. For example, if the `<blockstart>` value is 13200 and the block size configured for the multi-user system is 300, the actual source port range allocated to the user is 13200 through 13499. Each connection initiated by the user should use a unique source port number within the allocated range, enabling the firewall to identify the user based on its IP address-port-user mappings for enforcement of user- and group-based security rules. When a user exhausts all the ports allocated, the terminal server must send a new `<login>` message allocating a new port range for the user so that the firewall can update the IP address-port-user mapping. In addition, a single username can

 © Palo Alto Networks, Inc.

have multiple blocks of ports mapped simultaneously. When the firewall receives a <logout> message that includes a <blockstart> parameter, it removes the corresponding IP address-port-user mapping from its mapping table. When the firewall receives a <logout> message with a username and IP address, but no <blockstart>, it removes the user from its table. And, if the firewall receives a <logout> message with an IP address only, it removes the multi-user system and all mappings associated with it.

The XML files that the terminal server sends to the firewall can contain multiple message types and the messages do not need to be in any particular order within the file. However, upon receiving an XML file that contains multiple message types, the firewall will process them in the following order: multiusersystem requests first, followed by logins, then logouts.

The following workflow provides an example of how to use the User-ID XML API to send user mappings from a non-Windows terminal server to the firewall.

**Use the User-ID XML API to Map Non-Windows Terminal Services Users**

Step 1 — Generate the API key that will be used to authenticate the API communication between the firewall and the Terminal server. To generate the key you must provide login credentials for an administrative account; the API is available to all administrators (including role-based administrators with XML API privileges enabled).

Any special characters in the password must be URL/percent-encoded.

From a browser, log in to the firewall. Then, to generate the API key for the firewall, open a new browser window and enter the following URL:

```
https://<Firewall-IPaddress>/api/?type=keygen&user=<username>&
password=<password>
```

Where <Firewall-IPaddress> is the IP address or FQDN of the firewall and <username> and <password> are the credentials for the administrative user account on the firewall. For example:

```
https://10.1.2.5/api/?type=keygen&user=admin&password=admin
```

The firewall responds with a message containing the key, for example:

```
<response status="success">
   <result>
      <key>k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg=</key>
   </result>
</response>
```

**Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)**

Step 2 (Optional) Generate a setup message that the terminal server will send to specify the port range and block size of ports per user that your terminal services agent uses.

If the terminal services agent does not send a setup message, the firewall will automatically create a terminal server agent configuration using the following default settings upon receipt of the first login message:

- Default port range: 1025 to 65534
- Per user block size: 200
- Maximum number of multi-user systems: 1000

The following shows a sample setup message:

```
<uid-message>
  <payload>
    <multiusersystem>
      <entry ip="10.1.1.23" startport="20000"
        endport="39999" blocksize="100">
    </multiusersystem>
  </payload>
  <type>update</type>
  <version>1.0</version>
</uid-message>
```

where `entry ip` specifies the IP address assigned to terminal server users, `startport` and `endport` specify the port range to use when assigning ports to individual users, and `blocksize` specifies the number of ports to assign to each user. The maximum blocksize is 4000 and each multi-user system can allocate a maximum of 1000 blocks.

If you define a custom blocksize and or port range, keep in mind that you must configure the values such that every port in the range gets allocated and that there are no gaps or unused ports. For example, if you set the port range to 1000-1499, you could set the block size to 100, but not to 200. This is because if you set it to 200, there would be unused ports at the end of the range.

Step 3 Create a script that will extract the login events and create the XML input file to send to the firewall.

Make sure the script enforces assignment of port number ranges at fixed boundaries with no port overlaps. For example, if the port range is 1000-1999 and the block size is 200, acceptable blockstart values would be 1000, 1200, 1400, 1600, or 1800. Blockstart values of 1001, 1300, or 1850 would be unacceptable because some of the port numbers in the range would be left unused.

The login event payload that the terminal server sends to the firewall can contain multiple login events.

The following shows the input file format for a user-ID XML login event:

```
<uid-message>
<payload>
<login>
<entry name="acme\jjaso" ip="10.1.1.23" blockstart="20000">
<entry name="acme\jparker" ip="10.1.1.23" blockstart="20100">
<entry name="acme\ccrisp" ip="10.1.1.23" blockstart="21000">
</login>
</payload>
<type>update</type>
<version>1.0</version>
</uid-message>
```

The firewall uses this information to populate its user mapping table. Based on the mappings extracted from the example above, if the firewall received a packet with a source address and port of 10.1.1.23:20101, it would map the request to user jparker for policy enforcement.

Each multi-user system can allocate a maximum of 1000 port blocks.

 © Palo Alto Networks, Inc.

**Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)**

| | | |
|---|---|---|
| Step 4 | Create a script that will extract the logout events and create the XML input file to send to the firewall.<br><br>Upon receipt of a `logout` event message with a `blockstart` parameter, the firewall removes the corresponding IP address-port-user mapping. If the `logout` message contains a username and IP address, but no `blockstart` parameter, the firewall removes all mappings for the user. If the `logout` message contains an IP address only, the firewall removes the multi-user system and all associated mappings. | The following shows the input file format for a User-ID XML logout event:<br>`<uid-message>`<br>`  <payload>`<br>`  <logout>`<br>`  <entry name="acme\jjaso" ip="10.1.1.23" blockstart="20000">`<br>`  <entry name="acme\ccrisp" ip="10.1.1.23">`<br>`  <entry ip="10.2.5.4">`<br>`  </logout>`<br>`  </payload>`<br>`  <type>update</type>`<br>`  <version>1.0</version>`<br>`</uid-message>`<br><br>You can also clear the multiuser system entry from the firewall using the following CLI command: **`clear xml-api multiusersystem`** |
| Step 5 | Make sure that the scripts you create include a way to dynamically enforce that the port block range allocated using the XML API matches the actual source port assigned to the user on the terminal server and that the mapping is removed when the user logs out or the port allocation changes. | One way to do this would be to use netfilter NAT rules to hide user sessions behind the specific port ranges allocated via the XML API based on the uid. For example, to ensure that a user with the user ID jjaso is mapped to a source network address translation (SNAT) value of 10.1.1.23:20000-20099, the script you create should include the following:<br>`[root@ts1 ~]# iptables -t nat -A POSTROUTING -m owner --uid-owner jjaso -p tcp -j SNAT --to-source 10.1.1.23:20000-20099`<br>Similarly, the scripts you create should also ensure that the IP table routing configuration dynamically removes the SNAT mapping when the user logs out or the port allocation changes:<br>`[root@ts1 ~]# iptables -t nat -D POSTROUTING 1` |
| Step 6 | Define how to package the XML input files containing the setup, login, and logout events into wget or cURL messages for transmission to the firewall. | To apply the files to the firewall using wget:<br>`> wget --post file <filename> "https://<Firewall-IPaddress>/api/?type=user-id&key=<key>&file-name=<input_filename.xml>&client=wget&vsys=<VSYS_name>"`<br>For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key `k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg` using wget would look as follows:<br>`> wget --post file login.xml "https://10.2.5.11/api/?type=user-id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg&file-name=login.xml&client=wget&vsys=vsys1"`<br>To apply the file to the firewall using cURL:<br>`> curl --form file=@<filename> https://<Firewall-IPaddress>/api/?type=user-id&key=<key>&vsys=<VSYS_name>`<br>For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key `k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg` using cURL would look as follows:<br>`> curl --form file@login.xml "https://10.2.5.11/api/?type=user-id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg&vsys=vsys1"` |

| Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued) | |
|---|---|
| Step 7 Verify that the firewall is successfully receiving login events from the terminal servers. | Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:<br>To verify if the terminal server is connecting to the firewall over XML:<br>`admin@PA-5050> show user xml-api multiusersystem`<br>`Host Vsys Users Blocks`<br>`----------`<br>`10.5.204.43 vsys1 5 2`<br>To verify that the firewall is receiving mappings from a terminal server over XML:<br>`admin@PA-5050> show user ip-port-user-mapping all`<br>`Global max host index 1, host hash count 1`<br>`XML API Multi-user System 10.5.204.43`<br>`Vsys 1, Flag 3`<br>`Port range: 20000 - 39999`<br>`Port size: start 200; max 2000`<br>`Block count 100, port count 20000`<br>`20000-20199: acme\administrator`<br>`Total host: 1` |

 © Palo Alto Networks, Inc.

## Send User Mappings to User-ID Using the XML API

Although the User-ID functionality provides many out-of-the box methods for obtaining user mapping information, you may have some applications or devices that capture user information that cannot be natively integrated with User-ID. In this case you can use the User-ID XML API to create custom scripts that allow you to leverage existing user data and send it to the User-ID agent or directly to the firewall.

The User-ID XML API is a RESTful API that uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services. To leverage user data from an existing system—such as a custom application developed internally or another device that is not supported by one of the existing user mapping mechanisms—you can create custom scripts to extract the data and send it to the firewall or the User-ID agent using the XML API.

To enable an external system to send user mapping information to the User-ID agent or directly to the firewall, you can create scripts that extract the user login and logout events and use them for input to the User-ID XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget using the firewall's API key for secure communication. For more details, refer to the PAN-OS XML API Usage Guide.

# Configure a Firewall to Share User Mapping Data with Other Firewalls

Because policy is local to each firewall, each firewall needs current user mapping and group mapping information to accurately enforce policy by user and group. To simplify administration, you can configure one firewall to be the *redistribution firewall* that collects all the mapping information and shares it with other firewalls. The redistribution firewall can share only the information it collects using local methods (for example, the PAN-OS integrated User-ID agent or Captive Portal); it can't share the information collected from the Windows-based User-ID and Terminal Services agents. You configure the receiving firewalls to retrieve the mapping information from the redistribution firewall; they don't need to communicate directly with domain servers.

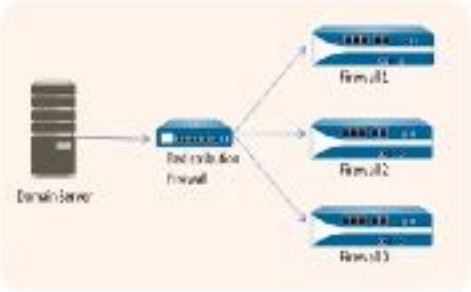

The following procedure describes how to set up redistribution of User-ID information.

| Configure a Firewall to Share User Mapping Data with Other Firewalls | | |
|---|---|---|
| Step 1 | Configure the redistribution firewall.<br><br>User-ID configurations apply to a single virtual system only. To redistribute User-ID mappings from multiple virtual systems, you must configure the user mapping settings on each virtual system separately, using a unique pre-shared key in each configuration. | 1. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User-ID Agent Setup section.<br>2. Select **Redistribution**.<br>3. Enter a **Collector Name**.<br>4. Enter and confirm the **Pre-Shared Key** that will enable other firewalls to connect to this firewall to retrieve user mapping information.<br>5. Click **OK** to save the redistribution configuration. |
| Step 2 | Create an interface management profile that enables the User-ID service and attach it to the interface that the other firewalls will connect to in order to retrieve user mappings. | 1. Select **Network > Network Profiles > Interface Mgmt** and click **Add**.<br>2. Enter a **Name** for the profile and then select the Permitted Services. At a minimum, select **User-ID Service** and **HTTPS**.<br>3. Click **OK** to save the profile.<br>4. Select **Network > Interfaces > Ethernet** and select the interface you plan to use for redistribution.<br>5. On the **Advanced > Other Info** tab, select the **Management Profile** you just created.<br>6. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

| Configure a Firewall to Share User Mapping Data with Other Firewalls (Continued) | | |
|---|---|---|
| Step 3 | Configure the other firewalls to retrieve user mappings from the redistribution firewall.<br> If the redistribution firewall has multiple virtual systems configured for redistribution, make sure you are using the pre-shared key that corresponds to the virtual system from which you want this firewall to retrieve User-ID mappings. | Perform the following steps on each firewall that you want to be able to retrieve user mappings:<br>1. Select **Device > User Identification > User-ID Agents**.<br>2. Click **Add** and enter a User-ID agent **Name** for the redistribution firewall.<br>3. Enter the hostname or IP address of the firewall interface that you configured for redistribution in the **Host** field.<br>4. Enter 5007 as the **Port** number on which the redistribution firewall will listen for User-ID requests.<br>5. Enter the **Collector Name** that you specified in the redistribution firewall configuration (Step 1-3).<br>6. Enter and confirm the **Collector Pre-Shared Key**. The key value you enter here must match the value configured on the redistribution firewall (Step 1-4).<br>7. (Optional) If you are using the redistribution firewall to retrieve group mappings in addition to user mappings, select the **Use as LDAP Proxy** check box.<br>8. (Optional) If you are using the redistribution firewall for Captive Portal authentication, select the **Use for NTLM Authentication** check box.<br>9. Make sure that the configuration is **Enabled** and then click **OK**.<br>10. **Commit** the changes. |
| Step 4 | Verify the configuration. | On the **User-ID Agents** tab, verify that the redistribution firewall entry you just added shows a green icon in the **Connected** column.<br>If a red icon appears, check traffic logs (**Monitor > Logs > Traffic**) to identify the issue. You can also check to see if any user mapping data has been received by running the following operational commands from the CLI:<br>`show user ip-user-mapping` (to view user mapping information on the dataplane)<br>`show user ip-user-mapping-mp` (to view mappings on the management plane). |

# Enable User- and Group-Based Policy

To enable security policy based on users and user groups, you must enable User-ID for each zone that contains users you want to identify. You can then define policy rules that allow or deny traffic based on username or group membership. Additionally, you can create Captive Portal rules to enable identification for IP addresses that don't yet have any user data associated with them.

For users with multiple usernames, see Enable Policy for Users with Multiple Accounts.

| Enable User- and Group-Based Policy | | |
|---|---|---|
| Step 1 | Enable User-ID on the source zones that contain the users who will send requests that require user-based access controls. | 1. Select **Network > Zones** and click the Name of the zone.<br>2. Select the **Enable User Identification** check box and click **OK**. |
| Step 2 | (Optional) Configure the firewall to read the IP addresses of users from the X-Forwarded-For (XFF) header in client requests for web services when the firewall is between the Internet and a proxy server that would otherwise hide the user IP addresses.<br><br>The firewall matches the IP addresses with usernames that your policy rules reference so that those rules can control and log access for the associated users and groups. For details, see Identify Users Connected through a Proxy Server. | 1. Select **Device > Setup > Content-ID** and edit the X-Forwarded-For Headers settings.<br>2. Select the **X-Forwarded-For Header in User-ID** check box.<br> Selecting the **Strip-X-Forwarded-For Header** check box doesn't disable the use of XFF headers for user attribution in policy rules; the firewall zeroes out the XFF value only after using it for user attribution.<br>3. Click **OK** to save your changes. |

 © Palo Alto Networks, Inc.

| Enable User- and Group-Based Policy (Continued) | |
|---|---|
| Step 3 Create security rules based on user and user group.<br><br> As a best practice, create rules based on group rather than user whenever possible. This prevents you from having to continually update your rules (which requires a commit) whenever your user base changes. | 1. After configuring User-ID, you will be able to choose a user name or group name when defining the source or destination of a security rule:<br>a. Select **Policies > Security** and click **Add** to create a new rule or click on an existing rule name to open the Security Policy Rule dialog.<br>b. Specify which users and/or groups to match in the rule in one of the following ways:<br>– If you want to specify specific users/groups as matching criteria, select the **User** tab and click the **Add** button in the Source User section to display a list of users and groups discovered by the firewall group mapping function. Select the users and/or groups to add to the rule.<br>– If you want the rule to match any user who has or has not successfully authenticated and you don't need to know the specific user or group name, select **known-user** or **unknown** from the drop-down above the **Source User** list.<br>2. Configure the rest of the rule as appropriate and then click **OK** to save it. For details on other fields in the security rule, see Set Up Basic Security Policies. |
| Step 4 Create your Captive Portal rules. | 1. Select **Policies > Captive Portal**.<br>2. Click **Add** and enter a **Name** for the rule.<br>3. Define the matching criteria for the rule by completing the **Source**, **Destination**, and **Service/URL Category** tabs as appropriate to match the traffic you want to authenticate. The matching criteria on these tabs is the same as the criteria you define when creating a security rule. See Set Up Basic Security Policies for details.<br>4. Define the **Action** to take on traffic that matches the rule:<br>• **no-captive-portal**—Allow traffic to pass without presenting a Captive Portal page for authentication.<br>• **web-form**—Present a Captive Portal page for the user to explicitly enter authentication credentials or use client certificate authentication.<br>• **browser-challenge**—Transparently obtain user authentication credentials. If you select this action, you must enable Kerberos Single Sign-On (SSO) or NT LAN Manager (NTLM) authentication when you Configure Captive Portal. If Kerberos SSO authentication fails, the firewall falls back to NTLM authentication. If you didn't configure NTLM, or NTLM authentication fails, the firewall falls back to **web-form** authentication.<br>5. Click **OK** and **Commit**. |

# Enable Policy for Users with Multiple Accounts

If a user in your organization has multiple responsibilities, that user might have multiple usernames (accounts), each with distinct privileges for accessing a particular set of services, but with all the usernames sharing the same IP address (the client system of the user). However, the User-ID agent can map any one IP address (or IP address and port range for terminal server users) to only one username for enforcing policy, and you can't predict which username the agent will map. To control access for all the usernames of a user, you must make adjustments to the rules, user groups, and User-ID agent.

For example, say the firewall has a rule that allows username corp_user to access email and a rule that allows username admin_user to access a MySQL server. The user logs in with either username from the same client IP address. If the User-ID agent maps the IP address to corp_user, then whether the user logs in as corp_user or admin_user, the firewall identifies that user as corp_user and allows access to email but not the MySQL server. On the other hand, if the User-ID agent maps the IP address to admin_user, the firewall always identifies the user as admin_user regardless of login and allows access to the MySQL server but not email. The following steps describe how to enforce both rules in this example.

| **Enable Policy for a User with Multiple Accounts** | | |
|---|---|---|
| Step 1 | Configure a user group for each service that requires distinct access privileges.<br><br>In this example, each group is for a single service (email or MySQL server). However, it is common to configure each group for a set of services that require the same privileges (for example, one group for all basic user services and one group for all administrative services). | If your organization already has user groups that can access the services that the user requires, simply add the username that is used for less restricted services to those groups. In this example, the email server requires less restricted access than the MySQL server, and corp_user is the username for accessing email. Therefore, you add corp_user to a group that can access email (corp_employees) and to a group that can access the MySQL server (network_services).<br><br>If adding a username to a particular existing group would violate your organizational practices, you can create a custom group based on an LDAP filter. For this example, say network_services is a custom group, which you configure as follows:<br>1. Select **Device > User Identification > Group Mapping Settings** and **Add** a group mapping configuration with a unique **Name**.<br>2. Select an LDAP **Server Profile** and ensure the **Enabled** check box is enabled.<br>3. Select the **Custom Group** tab and **Add** a custom group with network_services as a **Name**.<br>4. Specify an **LDAP Filter** that matches an LDAP attribute of corp_user and click **OK**.<br>5. Click **OK** and **Commit**.<br><br>Later, if other users that are in the group for less restricted services are given additional usernames that access more restricted services, you can add those usernames to the group for more restricted services. This scenario is more common than the inverse; a user with access to more restricted services usually already has access to less restricted services. |

 © Palo Alto Networks, Inc.

| Enable Policy for a User with Multiple Accounts (Continued) | | |
|---|---|---|
| Step 2 | Configure the rules that control user access based on the groups you just configured. | Enable User- and Group-Based Policy:<br>1. Configure a security rule that allows the corp_employees group to access email.<br>2. Configure a security rule that allows the network_services group to access the MySQL server. |
| Step 3 | Configure the ignore list of the User-ID agent.<br>This ensures that the User-ID agent maps the client IP address only to the username that is a member of the groups assigned to the rules you just configured. The ignore list must contain all the usernames of the user that are not members of those groups. | In this example, you add admin_user to the ignore list of the Windows-based User-ID agent to ensure that it maps the client IP address to corp_user. This guarantees that, whether the user logs in as corp_user or admin_user, the firewall identifies the user as corp_user and applies both rules that you configured because corp_user is a member of the groups that the rules reference.<br>1. Create an `ignore_user_list.txt` file.<br>2. Open the file and add admin_user.<br>If you later add more usernames, each must be on a separate line.<br>3. Save the file to the User-ID agent folder on the domain server where the agent is installed.<br>If you use the PAN-OS integrated User-ID agent, perform Step 6 under Configure User Mapping Using the PAN-OS Integrated User-ID Agent to configure the ignore list. |
| Step 4 | Configure endpoint authentication for the restricted services.<br>This enables the endpoint to verify the credentials of the user and preserves the ability to enable access for users with multiple usernames. | In this example, you have configured a firewall rule that allows corp_user, as a member of the network_services group, to send a service request to the MySQL server. You must now configure the MySQL server to respond to any unauthorized username (such as corp_user) by prompting the user to enter the login credentials of an authorized username (admin_user).<br>If the user logs in to the network as admin_user, the user can then access the MySQL server without it prompting for the admin_user credentials again.<br>In this example, both corp_user and admin_user have email accounts, so the email server won't prompt for additional credentials regardless of which username the user entered when logging in to the network.<br>The firewall is now ready to enforce rules for a user with multiple usernames. |

# Verify the User-ID Configuration

After you configure group mapping and user mapping and enable User-ID on your security rules and Captive Portal rules, you should verify that it is working properly.

| Verify the User-ID Configuration | |
|---|---|
| Step 1 Verify that group mapping is working. | From the CLI, enter the following command:<br>**show user group-mapping statistics** |
| Step 2 Verify that user mapping is working. | If you are using the on-device User-ID agent, you can verify this from the CLI using the following command:<br>**show user ip-user-mapping-mp all**<br>`IP Vsys From User Timeout (sec)`<br>`----------`<br>`192.168.201.1 vsys1 UIA acme\george 210`<br>`192.168.201.11 vsys1 UIA acme\duane 210`<br>`192.168.201.50 vsys1 UIA acme\betsy 210`<br>`192.168.201.10 vsys1 UIA acme\administrator 210`<br>`192.168.201.100 vsys1 AD acme\administrator 748`<br>`Total: 5 users`<br>`*: WMI probe succeeded` |
| Step 3 Test your security rule. | • From a machine in the zone where User-ID is enabled, attempt to access sites and applications to test the rules you defined in your policy and ensure that traffic is allowed and denied as expected.<br>• You can also use the `test security-policy-match` command to determine whether the policy is configured correctly. For example, suppose you have a rule that blocks user duane from playing World of Warcraft; you could test the policy as follows:<br>`test security-policy-match application worldofwarcraft source-user acme\duane source any destination any destination-port any protocol 6`<br>`"deny worldofwarcraft" {`<br>`from corporate;`<br>`source any;`<br>`source-region any;`<br>`to internet;`<br>`destination any;`<br>`destination-region any;`<br>`user acme\duane;`<br>`category any;`<br>`application/service worldofwarcraft;`<br>`action deny;`<br>`terminal no;`<br>`}` |

 © Palo Alto Networks, Inc.

**Verify the User-ID Configuration (Continued)**

Step 4 Test your Captive Portal configuration.

1. From the same zone, go to a machine that is not a member of your directory, such as a Mac OS system, and try to ping to a system external to the zone. The ping should work without requiring authentication.
2. From the same machine, open a browser and navigate to a web site in a destination zone that matches a Captive Portal rule you defined. You should see the Captive Portal web form.

3. Log in using the correct credentials and confirm that you are redirected to the requested page.
4. You can also test your Captive Portal policy using the `test cp-policy-match` command as follows:

```
test cp-policy-match from corporate to internet
source 192.168.201.10 destination 8.8.8.8
Matched rule: 'captive portal' action: web-form
```

Step 5 Verify that user names are displayed in the log files (**Monitor > Logs**).

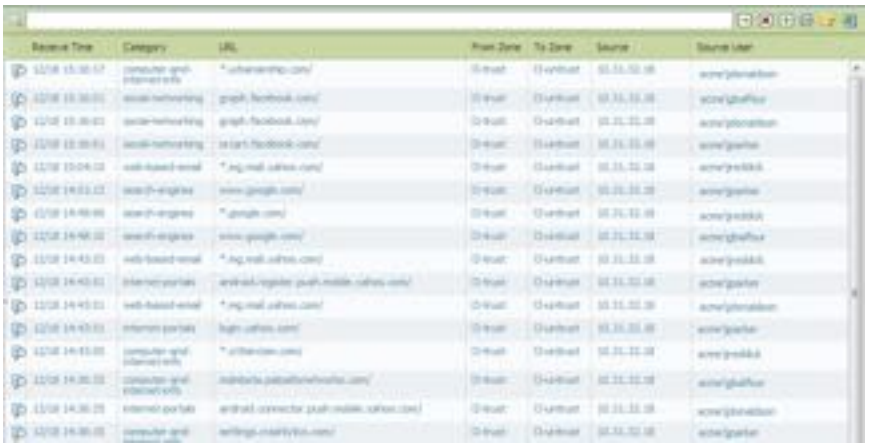

**Verify the User-ID Configuration (Continued)**

Step 6 Verify that user names are displayed in reports (**Monitor > Reports**). For example, when drilling down into the denied applications report, you should see a list of the users who attempted to access the applications as in the following example.

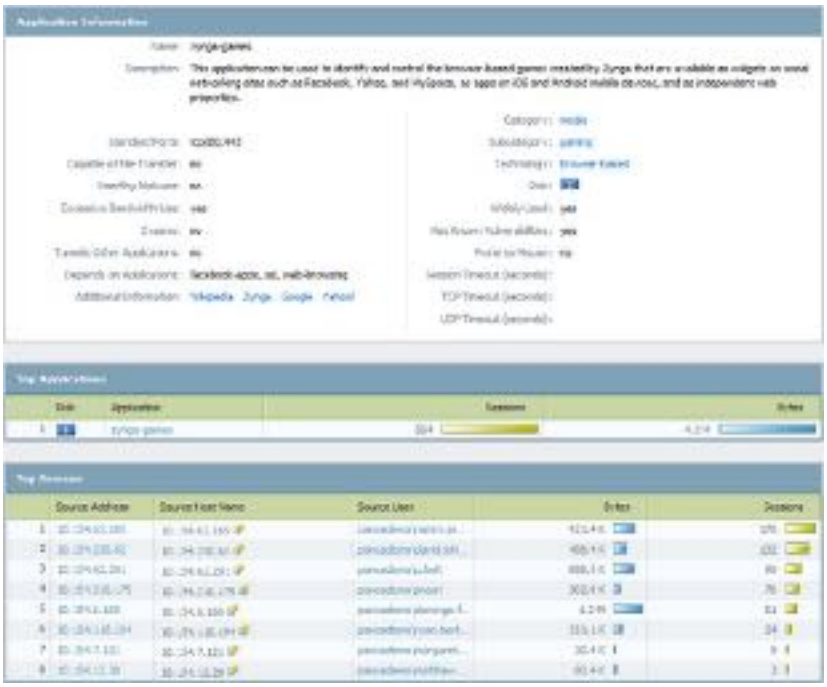

 © Palo Alto Networks, Inc.

# Deploy User-ID in a Large-Scale Network

The following topics describe a User-ID deployment in which you use Windows Log Forwarding and Global Catalog servers to simplify user and group mapping in a large-scale network of Microsoft Active Directory (AD) domain controllers or Exchange servers:

- Windows Log Forwarding and Global Catalog Servers
- Plan Your User-ID Implementation for a Large-Scale Network
- Configure Windows Log Forwarding
- Configure User-ID for a Large-Scale Network

## Windows Log Forwarding and Global Catalog Servers

Because each User-ID agent can monitor up to 100 servers, you need multiple User-ID agents to monitor a network with hundreds of AD domain controllers or Exchange servers. Creating and managing numerous User-ID agents involves considerable administrative overhead, especially in expanding networks where tracking new domain controllers is difficult. Windows Log Forwarding enables you to minimize the administrative overhead by reducing the number of servers to monitor and thereby reducing the number of User-ID agents to manage. When you configure Windows Log Forwarding, multiple domain controllers export their login events to a single domain member from which a User-ID agent collects the user mapping information.

You can configure Windows Log Forwarding for Windows Server versions 2003, 2008, 2008 R2, 2012, and 2012 R2. Windows Log Forwarding is not available for non-Microsoft servers.

To collect username to group mapping information in a large-scale network, you can configure the firewall to query a Global Catalog server that receives account information from the domain controllers.

The following figure illustrates user mapping and group mapping for a large-scale network in which the firewall uses a Windows-based User-ID agent. See Plan Your User-ID Implementation for a Large-Scale Network to determine if this deployment suits your network.

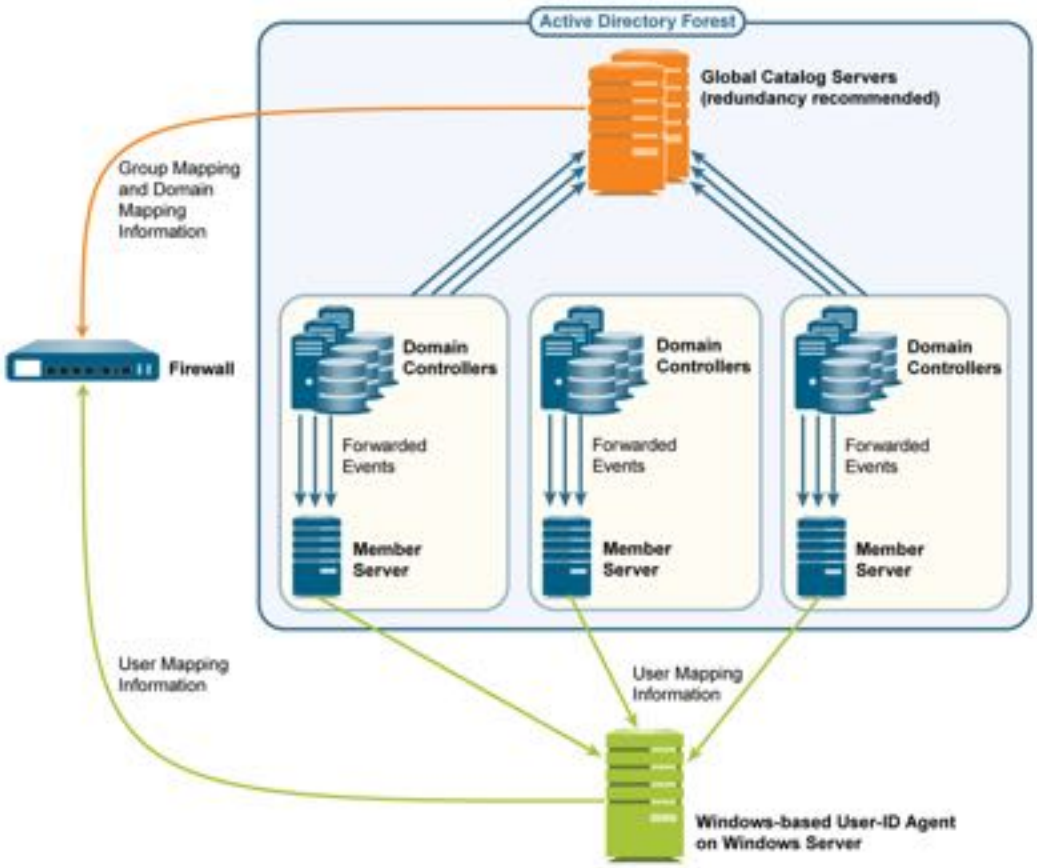


## Plan Your User-ID Implementation for a Large-Scale Network

When deciding whether to use Windows Log Forwarding and Global Catalog servers for your User-ID implementation, consult your system administrator to determine:

- Bandwidth required for domain controllers to forward login events to member servers. The bandwidth is a multiple of the login rate (number of logins per minute) of the domain controllers and the byte size of each login event.

  Note that domain controllers won't forward their entire security logs; they forward only the events that the user mapping process requires per login: three events for Windows Server 2003 or four events for Windows Server 2008/2012 and MS Exchange.

- Whether the following network elements support the required bandwidth:
  - Domain controllers—They must support the processing load associated with forwarding the events.
  - Member Servers—They must support the processing load associated with receiving the events.
  - Connections—The geographic distribution (local or remote) of the domain controllers, member servers, and Global Catalog servers is a factor. Generally, a remote distribution supports less bandwidth.

 © Palo Alto Networks, Inc.

## Configure Windows Log Forwarding

To configure Windows Log Forwarding, you need administrative privileges for configuring group policies on Windows servers. Configure Windows Log Forwarding on every member server that will collect login events from domain controllers. The following is an overview of the tasks; consult your Windows Server documentation for the specific steps.

| Configure Windows Log Forwarding | |
|---|---|
| Step 1 | On every member server that will collect security events, enable event collection, add the domain controllers as event sources, and configure the event collection query (subscription). The events you specify in the subscription vary by domain controller platform:<br>• Windows Server 2003—The event IDs for the required events are 672 (Authentication Ticket Granted), 673 (Service Ticket Granted), and 674 (Ticket Granted Renewed).<br>• Windows Server 2008/2012 (including R2) or MS Exchange—The event IDs for the required events are 4768 (Authentication Ticket Granted), 4769 (Service Ticket Granted), 4770 (Ticket Granted Renewed), and 4624 (Logon Success).<br> You must forward events to the security logs location on the member servers, not to the default forwarded logs location.<br> To forward events as quickly as possible, select the **Minimize Latency** option when configuring the subscription. |
| Step 2 | Configure a group policy to enable Windows Remote Management (WinRM) on the domain controllers. |
| Step 3 | Configure a group policy to enable Windows Event Forwarding on the domain controllers. |

## Configure User-ID for a Large-Scale Network

| Configure User-ID for a Large-Scale Network | | |
|---|---|---|
| Step 1 | Configure Windows Log Forwarding on the member servers that will collect login events. | Configure Windows Log Forwarding. This step requires administrative privileges for configuring group policies on Windows servers. |
| Step 2 | Install the Windows-based User-ID agent. | Install the User-ID Agent on a Windows server that can access the member servers. The Windows server can be inside or outside the Active Directory forest; it doesn't need to be a member server itself. |

| Configure User-ID for a Large-Scale Network (Continued) | | |
|---|---|---|
| Step 3 | Configure the User-ID agent to collect user mapping information from the member servers. | 1. Start the Windows-based User-ID agent.<br>2. Select **User Identification > Discovery** and perform the following steps for each member server that will receive events from domain controllers:<br>a. In the Servers section, click **Add** and enter a **Name** to identify the member server.<br>b. In the **Server Address** field, enter the FQDN or IP address of the member server.<br>c. For the **Server Type**, select **Microsoft Active Directory**.<br>d. Click **OK** to save the server entry.<br>3. Configure the remaining User-ID agent settings: see Configure the User-ID Agent for User Mapping. |
| Step 4 | Configure an LDAP server profile to specify how the firewall connects to the Global Catalog servers (up to four) for group mapping information.<br>To improve availability, use at least two Global Catalog servers for redundancy.<br>You can collect group mapping information only for universal groups, not local domain groups (subdomains). | 1. Select **Device > Server Profiles > LDAP**, click **Add**, and enter a **Name** for the profile.<br>2. In the Servers section, for each Global Catalog, click **Add** and enter the server **Name**, IP address (**LDAP Server**), and **Port**. For a plaintext or Start Transport Layer Security (Start TLS) connection, use **Port** 3268. For an LDAP over SSL connection, use **Port** 3269. If the connection will use Start TLS or LDAP over SSL, select the **Require SSL/TLS secured connection** check box.<br>3. In the **Base DN** field, enter the Distinguished Name (DN) of the point in the Global Catalog server where the firewall will start searching for group mapping information (for example, `DC=acbdomain,DC=com`).<br>4. For the **Type**, select **active-directory**.<br>5. Configure the remaining fields as necessary: see Add an LDAP server profile. |
| Step 5 | Configure an LDAP server profile to specify how the firewall connects to the servers (up to four) that contain domain mapping information.<br>User-ID uses this information to map DNS domain names to NetBIOS domain names. This mapping ensures consistent domain/username references in policy rules.<br>To improve availability, use at least two servers for redundancy. | The steps are the same as for the LDAP server profile you created for Global Catalogs in the Step 4, except for the following fields:<br>• **LDAP Server**—Enter the IP address of the domain controller that contains the domain mapping information.<br>• **Port**—For a plaintext or Start TLS connection, use **Port** 389. For an LDAP over SSL connection, use **Port** 636. If the connection will use Start TLS or LDAP over SSL, select the **Require SSL/TLS secured connection** check box.<br>• **Base DN**—Select the DN of the point in the domain controller where the firewall will start searching for domain mapping information. The value must start with the string: `cn=partitions,cn=configuration` (for example, `cn=partitions,cn=configuration,DC=acbdomain,DC=com`). |


© Palo Alto Networks, Inc.

| Configure User-ID for a Large-Scale Network (Continued) | | |
|---|---|---|
| Step 6 | Create a group mapping configuration for each LDAP server profile you created. | 1. Select **Device > User Identification > Group Mapping Settings**.<br>2. Click **Add** and enter a **Name** to identify the group mapping configuration.<br>3. Select the LDAP **Server Profile** and ensure the **Enabled** check box is selected.<br>4. Configure the remaining fields as necessary: see Map Users to Groups.<br> If the Global Catalog and domain mapping servers reference more groups than your security rules require, configure the **Group Include List** and/or **Custom Group** list to limit the groups for which User-ID performs mapping.<br>5. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# App-ID

To safely enable applications on your network, the Palo Alto Networks next-generation firewalls provide both an application and web perspective—App-ID and URL Filtering—to protect against a full spectrum of legal, regulatory, productivity, and resource utilization risks.

App-ID enables visibility into the applications on the network, so you can learn how they work and understand their behavioral characteristics and their relative risk. This application knowledge allows you to create and enforce security policy rules to enable, inspect, and shape desired applications and block unwanted applications. When you define policy rules to allow traffic, App-ID begins to classify traffic without any additional configuration.


- App-ID Overview
- Manage Custom or Unknown Applications
- Manage New App-IDs Introduced in Content Releases
- Use Application Objects in Policy
- Applications with Implicit Support
- Application Level Gateways
- Disable the SIP Application-level Gateway (ALG)


© Palo Alto Networks, Inc.

# App-ID Overview

App-ID, a patented traffic classification system only available in Palo Alto Networks firewalls, determines what an application is irrespective of port, protocol, encryption (SSH or SSL) or any other evasive tactic used by the application. It applies multiple classification mechanisms—application signatures, application protocol decoding, and heuristics—to your network traffic stream to accurately identify applications.

Here's how App-ID identifies applications traversing your network:

- Traffic is matched against policy to check whether it is allowed on the network.
- Signatures are then applied to allowed traffic to identify the application based on unique application properties and related transaction characteristics. The signature also determines if the application is being used on its default port or it is using a non-standard port. If the traffic is allowed by policy, the traffic is then scanned for threats and further analyzed for identifying the application more granularly.
- If App-ID determines that encryption (SSL or SSH) is in use, and a Decryption policy rule is in place, the session is decrypted and application signatures are applied again on the decrypted flow.
- Decoders for known protocols are then used to apply additional context-based signatures to detect other applications that may be tunneling inside of the protocol (for example, Yahoo! Instant Messenger used across HTTP). Decoders validate that the traffic conforms to the protocol specification and provide support for NAT traversal and opening dynamic pinholes for applications such as SIP and FTP.
- For applications that are particularly evasive and cannot be identified through advanced signature and protocol analysis, heuristics or behavioral analysis may be used to determine the identity of the application.

When the application is identified, the policy check determines how to treat the application, for example—block, or allow and scan for threats, inspect for unauthorized file transfer and data patterns, or shape using QoS.

 © Palo Alto Networks, Inc.

# Manage Custom or Unknown Applications

Palo Alto Networks provides weekly application updates to identify new App-ID signatures. By default, App-ID is always enabled on the firewall, and you don't need to enable a series of signatures to identify well-known applications. Typically, the only applications that are classified as *unknown* traffic—*tcp, udp or non-syn-tcp*—in the ACC and the traffic logs are commercially available applications that have not yet been added to App-ID, internal or custom applications on your network, or potential threats.

On occasion, the firewall may report an application as unknown for the following reasons:

- Incomplete data—A handshake took place, but no data packets were sent prior to the timeout.
- Insufficient data—A handshake took place followed by one or more data packets; however, not enough data packets were exchanged to identify the application.

The following choices are available to handle unknown applications:

- Create security policies to control unknown applications by unknown TCP, unknown UDP or by a combination of source zone, destination zone, and IP addresses.
- Request an App-ID from Palo Alto Networks—If you would like to inspect and control the applications that traverse your network, for any unknown traffic, you can record a packet capture. If the packet capture reveals that the application is a commercial application, you can submit this packet capture to Palo Alto Networks for App-ID development. If it is an internal application, you can create a custom App-ID and/or define an application override policy.
- Create a Custom Application with a signature and attach it to a security policy, or create a custom application and define an application override policy—A custom application allows you to customize the definition of the internal application—its characteristics, category and sub-category, risk, port, timeout—and exercise granular policy control in order to minimize the range of unidentified traffic on your network. Creating a custom application also allows you to correctly identify the application in the **ACC** and traffic logs and is useful in auditing/reporting on the applications on your network. For a custom application you can specify a signature and a pattern that uniquely identifies the application and attach it to a security policy that allows or denies the application.

  Alternatively, if you would like the firewall to process the custom application using fast path (Layer-4 inspection instead of using App-ID for Layer-7 inspection), you can reference the custom application in an application override policy rule. An application override with a custom application will prevent the session from being processed by the App-ID engine, which is a Layer-7 inspection. Instead it forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4, and thereby saves application processing time.

  For example, if you build a custom application that triggers on a host header *www.mywebsite.com*, the packets are first identified as *web-browsing* and then are matched as your custom application (whose parent application is web-browsing). Because the parent application is web-browsing, the custom application is inspected at Layer-7 and scanned for content and vulnerabilities.

  If you define an application override, the firewall stops processing at Layer-4. The custom application name is assigned to the session to help identify it in the logs, and the traffic is not scanned for threats.

For more details, refer to the following articles:

- Identifying Unknown Applications
- Video: How to Configure a Custom App-ID
- Custom Application Signatures

 © Palo Alto Networks, Inc.

# Manage New App-IDs Introduced in Content Releases

Installing new App-IDs included in a content release version can sometimes cause a change in policy enforcement for the now uniquely-identified application. Before installing a new content release, review the policy impact for new App-IDs and stage any necessary policy updates. Assess the treatment an application receives both before and after the new content is installed. You can then modify existing security policy rules using the new App-IDs contained in a downloaded content release (prior to installing the App-IDs). This enables you to simultaneously update your security policies and install new content, and allows for a seamless shift in policy enforcement. Alternatively, you can also choose to disable new App-IDs when installing a new content release version; this enables protection against the latest threats, while giving you the flexibility to enable the new App-IDs after you've had the chance to prepare any policy changes.

The following options enable you to assess the impact of new App-IDs on existing policy enforcement, disable (and enable) App-IDs, and seamlessly update policy rules to secure and enforce newly-identified applications:

- Review New App-IDs
- Disable or Enable App-IDs
- Prepare Policy Updates For Pending App-IDs

## Review New App-IDs

Review new App-ID signatures introduced in a Applications and/or Threats content update. For each new application signature introduced, you can preview the App-ID details, including a description of the application identified by the App-ID, other existing App-IDs that the new signature is dependent on (such as SSL or HTTP), and the category the application traffic received before the introduction of the new App-ID (for example, an application might be classified as web-browsing traffic before a App-ID signature is introduced that uniquely identifies the traffic). After reviewing the description and details for a new App-ID signature, review the App-ID signature impact on existing policy enforcement. When new application signatures are introduced, the newly-identified application traffic might no longer match to policies that previously enforced the application. Reviewing the policy impact for new application signatures enables you to identify the policies that will no longer enforce the application when the new App-ID is installed.

After downloading a new content release version, review the new App-IDs included in the content version and assess the impact of the new App-IDs on existing policy rules:

- Review New App-IDs Since Last Content Version
- Review New App-ID Impact on Existing Policy Rules

 © Palo Alto Networks, Inc.

# Review New App-IDs Since Last Content Version

**Review New App-IDs Available Since the Last Installed Content Release Version**

Step 1 Select **Device > Dynamic Updates** and select **Check Now** to refresh the list of available content updates.

Step 2 **Download** the latest Applications and Threats content update. When the content update is downloaded, an **Apps** link will appear in the Features column for that content update.

Step 3 Click the **Apps** link in the **Features** column to view details on newly-identified applications:

A list of App-IDs shows all new App-IDs introduced from the content version installed on the firewall, to the selected **Content Version**.

App-ID details that you can use to assess possible impact to policy enforcement include:

- **Depends on**—Lists the application signatures that this App-ID relies on to uniquely identify the application. If one of the application signatures listed in the **Depends On** field is disabled, the dependent App-ID is also disabled.
- **Previously Identified As**—Lists the App-IDs that matched to the application before the new App-ID was installed to uniquely identify the application.
- **App-ID Enabled**—All App-IDs display as enabled when a content release is downloaded, unless you choose to manually disable the App-ID signature before installing the content update (see Disable or Enable App-IDs).

  Multi-vsys firewalls display App-ID status as **vsys-specific**. This is because the status is not applied across virtual systems and must be individually enabled or disabled for each virtual system. To view the App-ID status for a specific virtual system, select **Objects > Applications**, select a **Virtual System**, and select the App-ID.

| Next Steps... | • Disable or Enable App-IDs.<br>• Prepare Policy Updates For Pending App-IDs. |
|---|---|

# Review New App-ID Impact on Existing Policy Rules

**Review the Impact of New App-ID Signatures on Existing Policy Rules**

Step 1 Select **Device > Dynamic Updates**.

Step 2 You can review the policy impact of new content release versions that are downloaded to the firewall. **Download** a new content release version, and click the **Review Policies** in the Action column. The **Policy review based on candidate configuration** dialog allows you to filter by **Content Version** and view App-IDs introduced in a specific release (you can also filter the policy impact of new App-IDs according to **Rulebase** and **Virtual System**).

Step 3 Select a new App-ID from the **Application** drop-down to view policy rules that currently enforce the application. The rules displayed are based on the applications signatures that match to the application before the new App-ID is installed (view application details to see the list of application signatures that an application was **Previously Identified As** before the new App-ID).

Step 4 Use the detail provided in the policy review to plan policy rule updates to take effect when the App-ID is installed and enabled to uniquely identify the application.

You can continue to Prepare Policy Updates For Pending App-IDs, or you can directly add the new App-ID to policy rules that the application was previously matched to by continuing to use the policy review dialog.

In the following example, the new App-ID adobe-cloud is introduced in a content release. Adobe-cloud traffic is currently identified as SSL and web-browsing traffic. Policy rules configured to enforce SSL or web-browsing traffic are listed to show what policy rules will be affected when the new App-ID is installed. In this example, the rule Allow SSL App currently enforces SSL traffic. To continue to allow adobe-cloud traffic when it is uniquely identified, and no longer identified as SSL traffic.

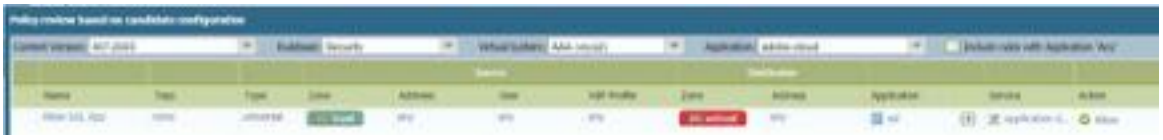

Add the new App-ID to existing policy rules, to allow the application traffic to continue to be enforced according to your existing security requirements when the App-ID is installed.

In this example, to continue to allow adobe-cloud traffic when it is uniquely identified by the new App-ID, and no longer identified as SSL traffic, add the new App-ID to the security policy rule Allow SSL App.

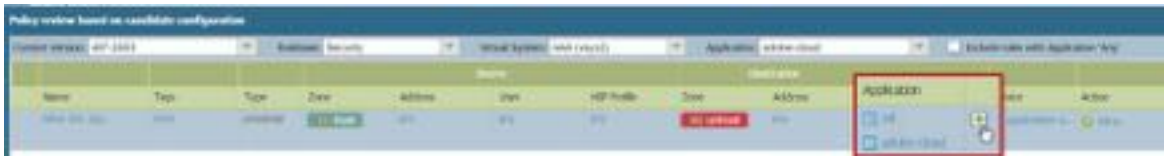

The policy rule updates take effect only when the application updates are installed.

| Next Steps... | • Disable or Enable App-IDs.<br>• Prepare Policy Updates For Pending App-IDs. |
|---|---|

 © Palo Alto Networks, Inc.

## Disable or Enable App-IDs

Disable new App-IDs included in a content release to immediately benefit from protection against the latest threats while continuing to have the flexibility to later enable App-IDs after preparing necessary policy updates. You can disable all App-IDs introduced in a content release, set scheduled content updates to automatically disable new App-IDs, or disable App-IDs for specific applications.

Policy rules referencing App-IDs only match to and enforce traffic based on enabled App-IDs.

Certain App-IDs cannot be disabled and only allow a status of enabled. App-IDs that cannot be disabled included some application signatures implicitly used by other App-IDs (such as unknown-tcp). Disabling a base App-ID could cause App-IDs which depend on the base App-ID to also be disabled. For example, disabling facebook-base will disable all other Facebook App-IDs.

| **Disable and Enable App-IDs** | |
|---|---|
| Disable all App-IDs in a content release or for scheduled content updates. | • To disable all new App-IDs introduced in a content release, select **Device > Dynamic Updates** and **Install** an Application and Threats content release. When prompted, select **Disable new apps in content update**. Select the check box to disable apps and continue installing the content update; this allows you to be protected against threats, and gives you the option to enable the apps at a later time.<br>• On the **Device > Dynamic Updates** page, select Schedule. Choose to **Disable new apps in content update** for downloads and installations of content releases. |
| Disable App-IDs for one application or multiple applications at a single time. | • To quickly disable a single application or multiple applications at the same time, click **Objects > Applications**. Select one or more application check box and click **Disable**.<br>• To review details for a single application, and then disable the App-ID for that application, select **Objects > Applications** and **Disable App-ID**. You can use this step to disable both pending App-IDs (where the content release including the App-ID is downloaded to the firewall but not installed) or installed App-IDs. |
| Enable App-IDs. | Enable App-IDs that you previously disabled by selecting **Objects > Applications**. Select one or more application check box and click **Enable** or open the details for a specific application and click **Enable App-ID**. |

## Prepare Policy Updates For Pending App-IDs

You can now stage seamless policy updates for new App-IDs. Release versions prior to PAN-OS 7.0 required you to install new App-IDs (as part of a content release) and then make necessary policy updates. This allowed for a period during which the newly-identified application traffic was not enforced, either by existing rules (that the traffic had matched to before being uniquely identified) or by rules that had yet to be created or modified to use the new App-ID.

*Pending* App-IDs can now be added to policy rules to prevent gaps in policy enforcement that could occur during the period between installing a content release and updating security policy. Pending App-IDs includes App-IDs that have been manually disabled, or App-IDs that are downloaded to the firewall but not installed. Pending App-IDs can be used to update policies both before and after installing a new content release. Though they can be added to policy rules, pending App-IDs are not enforced until the App-IDs are both installed and enabled on the firewall.

The names of App-IDs that have been manually disabled display as gray and italicized, to indicate the disabled status:

- Disabled App-ID listed on the **Objects > Applications** page:

- Disabled App-ID included in a security policy rule:

App-IDs that are included in a downloaded content release version might have an App-ID status of enabled, but App-IDs are not enforced until the corresponding content release version is installed.

 © Palo Alto Networks, Inc.

**Perform Seamless Policy Updates for New App-IDs**

To install the content release version now and then update policies:

Do this to benefit from new threat signatures immediately, while you review new application signatures and update your policies.

1. Select **Device > Dynamic Updates** and **Download** the latest content release version.
2. Review the Impact of New App-ID Signatures on Existing Policy Rules to assess the policy impact of new App-IDs.
3. **Install** the latest content release version. Before the content release is installed, you are prompted to **Disable new apps in content update**. Select the check box and continue to install the content release. Threat signatures included in the content release will be installed and effective, while new or updated App-IDs are disabled.
4. Select **Policies** and update **Security**, **QoS**, and **Policy Based Forwarding** rules to match to and enforce the now uniquely identified application traffic, using the pending App-IDs.
5. Select **Objects > Applications** and select one or multiple disabled App-IDs and click **Enable**.
6. **Commit** your changes to seamlessly update policy enforcement for new App-IDs.

To update policies now and then install the content release version:

1. Select **Device > Dynamic Updates** and **Download** the latest content release version.
2. Review the Impact of New App-ID Signatures on Existing Policy Rules to assess the policy impact of new App-IDs.
3. While reviewing the policy impact for new App-IDs, you can use the **Policy Review based on candidate configuration** to add a new App-ID to existing policy rules: .
4. The new App-ID is added to the existing rules as a disabled App-ID.
5. Continue to review the policy impact for all App-IDs included in the latest content release version by selecting App-IDs in the **Applications** drop-down. Add the new App-IDs to existing policies as needed. Click **OK** to save your changes.
6. **Install** the latest content release version.
7. **Commit** your changes to seamlessly update policy enforcement for new App-IDs.

© Palo Alto Networks, Inc.

# Use Application Objects in Policy


- Create an Application Group
- Create an Application Filter
- Create a Custom Application


## Create an Application Group

An application group is an object that contains applications that you want to treat similarly in policy. Application groups are useful for enabling access to applications that you explicitly sanction for use within your organization. Grouping sanctioned applications simplifies administration of your rulebases.: instead of having to update individual policy rules when there is a change in the applications you support, you can instead update only the affected application groups.

When deciding how to group applications, consider how you plan to enforce access to your sanctioned applications and create an application group that aligns with each of your policy goals. For example, you might have some applications that you will only allow your IT administrators to access, and other applications that you want to make available for any known user in your organization. In this case, you would create separate application groups for each of these policy goals. Although you generally want to enable access to applications on the default port only, you may want to group applications that are an exception to this and enforce access to those applications in a separate rule.

| **Create an Application Group** | |
|---|---|
| Step 1 | Select **Objects > Application Groups**. |
| Step 2 | **Add** a group and give it a descriptive **Name**. |
| Step 3 | (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall. |
| Step 4 | **Add** the applications you want in the group and then click **OK**.<br> |
| Step 5 | **Commit** the configuration. |


© Palo Alto Networks, Inc.

## Create an Application Filter

An application filter is an object that dynamically groups applications based on application attributes that you define, including category, subcategory, technology, risk factor, and characteristic. This is useful when you want to safely enable access to applications that you do not explicitly sanction, but that you want users to be able to access. For example, you may want to enable employees to choose their own office programs (such as Evernote, Google Docs, or Microsoft Office 365) for business use. To safely enable these types of applications, you could create an application filter that matches on the Category **business-systems** and the Subcategory **office-programs**. As new applications office programs emerge and new App-IDs get created, these new applications will automatically match the filter you defined; you will not have to make any additional changes to your policy rulebase to safely enable any application that matches the attributes you defined for the filter.

| **Create an Application Filter** | |
|---|---|
| Step 1 | Select **Objects > Application Filters**. |
| Step 2 | **Add** a filter and give it a descriptive **Name**. |
| Step 3 | (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall. |
| Step 4 | Define the filter by selecting attribute values from the Category, Subcategory, Technology, Risk, and Characteristic sections. As you select values, notice that the list of matching applications at the bottom of the dialog narrows. When you have adjusted the filter attributes to match the types of applications you want to safely enable, click **OK**. 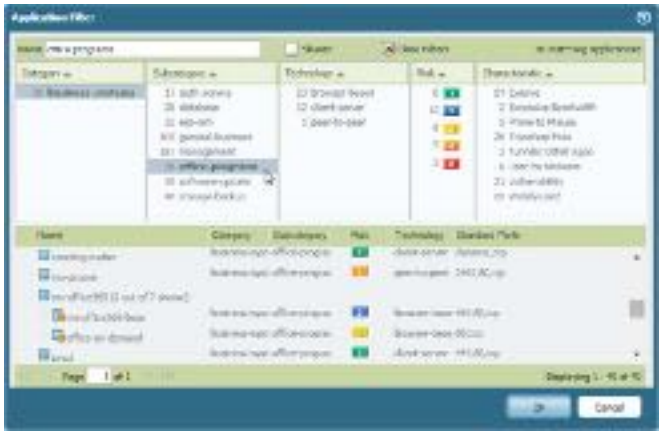 |
| Step 5 | **Commit** the configuration. |

## Create a Custom Application

To safely enable applications you must classify all traffic, across all ports, all the time. With App-ID, the only applications that are typically classified as unknown traffic—tcp, udp or non-syn-tcp—in the ACC and the Traffic logs are commercially available applications that have not yet been added to App-ID, internal or custom applications on your network, or potential threats.

If you are seeing unknown traffic for a commercial application that does not yet have an App-ID, you can submit a request for a new App-ID here: http://researchcenter.paloaltonetworks.com/submit-an-application/.

To ensure that your internal custom applications do not show up as unknown traffic, create a custom application. You can then exercise granular policy control over these applications in order to minimize the range of unidentified traffic on your network, thereby reducing the attack surface. Creating a custom application also allows you to correctly identify the application in the ACC and Traffic logs, which enables you to audit/report on the applications on your network.

To create a custom application, you must define the application attributes: its characteristics, category and sub-category, risk, port, timeout. In addition, you must define patterns or values that the firewall can use to match to the traffic flows themselves (the *signature*). Finally, you can attach the custom application to a security policy that allows or denies the application (or add it to an application group or match it to an application filter). You can also create custom applications to identify ephemeral applications with topical interest, such as ESPN3-Video for world cup soccer or March Madness.

In order to collect the right data to create a custom application signature, you'll need a good understanding of packet captures and how datagrams are formed. If the signature is created too broadly, you might inadvertently include other similar traffic; if it is defined too narrowly, the traffic will evade detection if it does not strictly match the pattern.

Custom applications are stored in a separate database on the firewall and this database is not impacted by the weekly App-ID updates.

The supported application protocol decoders that enable the firewall to detect applications that may be tunneling inside of the protocol include the following as of content update 424: HTTP, HTTPS, DNS, FTP, IMAP SMTP, Telnet, IRC (Internet Relay Chat), Oracle, RTMP, RTSP, SSH, GNU-Debugger, GIOP (Global Inter-ORB Protocol), Microsoft RPC, Microsoft SMB (also known as CIFS).

The following is a basic example of how to create a custom application.

 © Palo Alto Networks, Inc.

| Create a Custom Application | | |
|---|---|---|
| Step 1 | Gather information about the application that you will be able to use to write custom signatures.<br><br>To do this, you must have an understanding of the application and how you want to control access to it. For example, you may want to limit what operations users can perform within the application (such as uploading, downloading, or live streaming). Or you may want to allow the application, but enforce QoS policing. | • Capture application packets so that you can find unique characteristics about the application on which to base your custom application signature. One way to do this is to run a protocol analyzer, such as Wireshark, on the client system to capture the packets between the client and the server. Perform different actions in the application, such as uploading and downloading, so that you will be able to locate each type of session in the resulting packet captures (PCAPs).<br>• Because the firewall by default takes packet captures for all unknown traffic, if the firewall is between the client and the server you can view the packet capture for the unknown traffic directly from the Traffic log.<br>• Use the packet captures to find patterns or values in the packet *contexts* that you can use to create signatures that will uniquely match the application traffic. For example, look for string patterns in HTTP response or request headers, URI paths, or hostnames. For information on the different string contexts you can use to create application signatures and where you can find the corresponding values in the packet, refer to Creating Custom Threat Signatures. |
| Step 2 | Add the custom application. | 1. Select **Objects > Applications** and click **Add**.<br>2. On the **Configuration** tab, enter a **Name** and a **Description** for the custom application that will help other administrators understand why you created the application.<br>3. (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.<br>4. Define the application Properties and Characteristics.<br> 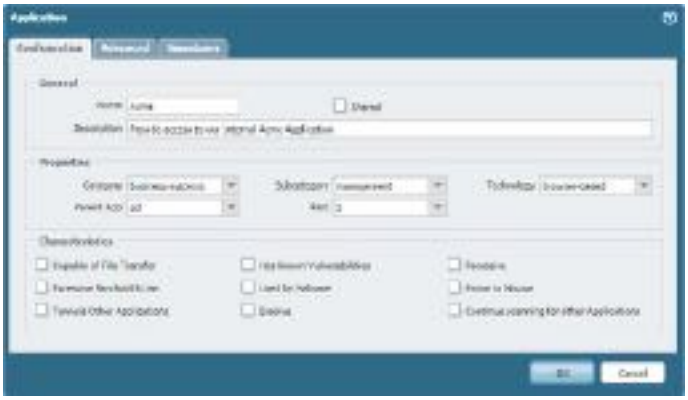 |

| Create a Custom Application (Continued) | |
|---|---|
| Step 3 Define details about the application, such as the underlying protocol, the port number the application runs on, the timeout values, and any types of scanning you want to be able to perform on the traffic. | On the **Advanced** tab, define settings that will allow the firewall to identify the application protocol:<br>• Specify the default ports or protocol that the application uses.<br>• Specify the session timeout values. If you don't specify timeout values, the default timeout values will be used.<br>• Indicate any type of additional scanning you plan to perform on the application traffic.<br>For example, to create a custom TCP-based application that runs over SSL, but uses port 4443 (instead of the default port for SSL, 443), you would specify the port number. By adding the port number for a custom application, you can create policy rules that use the default port for the application rather than opening up additional ports on the firewall. This improves your security posture.<br> 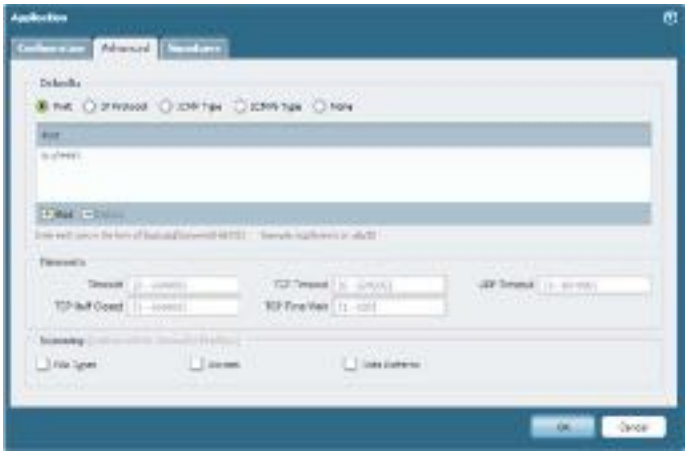 |

 © Palo Alto Networks, Inc.

**Create a Custom Application (Continued)**

Step 4 Define the criteria that the firewall will use to match the traffic to the new application.

You will use the information you gathered from the packet captures to specify unique string context values that the firewall can use to match patterns in the application traffic.

1. On the **Signatures** tab, click **Add** and define a **Signature Name** and optionally a **Comment** to provide information about how you intend to use this signature.
2. Specify the **Scope** of the signature: whether it matches to a full **Session** or a single **Transaction**.
3. Specify conditions to define signatures by clicking **Add And Condition** or **Add Or Condition**.
4. Select an **Operator** to define the type of match conditions you will use: **Pattern Match** or **Equal To**.
   - If you selected **Pattern Match**, select the **Context** and then use a regular expression to define the **Pattern** to match the selected context. Optionally, click **Add** to define a qualifier/value pair. The **Qualifier** list is specific to the **Context** you chose.
   - If you selected **Equal To**, select the **Context** and then use a regular expression to define the **Position** of the bytes in the packet header to use match the selected context. Choose from **first-4bytes** or **second-4bytes**. Define the 4-byte hex value for the **Mask** (for example, 0xffffff00) and **Value** (for example, 0xaabbccdd).

   For example, if you are creating a custom application for one of your internal applications, you could use the **ssl-rsp-certificate Context** to define a pattern match for the certificate response message of a SSL negotiation from the server and create a **Pattern** to match the commonName of the server in the message as shown here:

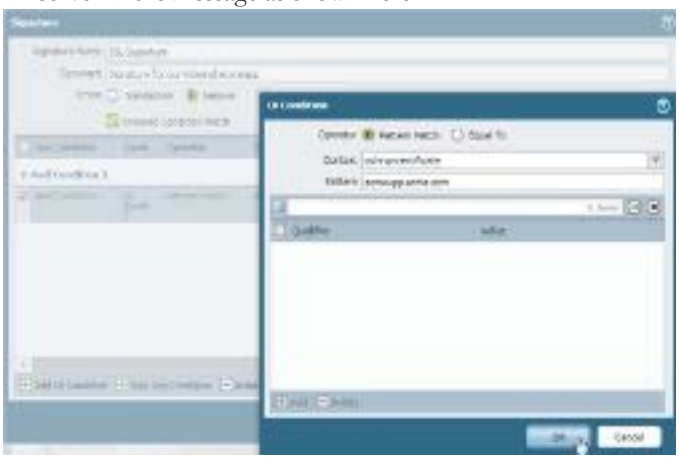

5. Repeat step 3 and 4 for each matching condition.
6. If the order in which the firewall attempts to match the signature definitions is important, make sure the **Ordered Condition Match** check box is selected and then order the conditions so that they are evaluated in the appropriate order. Select a condition or a group and click **Move Up** or **Move Down**. You cannot move conditions from one group to another.
7. Click **OK** to save the signature definition.

| Create a Custom Application (Continued) | | |
|---|---|---|
| Step 5 | Save the application. | 1. Click **OK** to save the custom application definition.<br>2. Click **Commit**. |
| Step 6 | Validate that traffic matches the custom application as expected. | 1. Select **Policies > Security** and **Add** a security policy rule to allow the new application.<br>2. Run the application from a client system that is between the firewall and the application and then check the Traffic logs (**Monitor > Traffic**) to make sure that you see traffic matching the new application (and that it is being handled per your policy rule). |

 © Palo Alto Networks, Inc.

# Applications with Implicit Support

When creating a policy to allow specific applications, you must also be sure that you are allowing any other applications on which the application depends. In many cases, you do not have to explicitly allow access to the dependent applications in order for the traffic to flow because the firewall is able to determine the dependencies and allow them implicitly. This implicit support also applies to custom applications that are based on HTTP, SSL, MS-RPC, or RTSP. Applications for which the firewall cannot determine dependent applications on time will require that you explicitly allow the dependent applications when defining your policies. You can determine application dependencies in Applipedia.

The following table lists the applications for which the firewall has implicit support (as of Content Update 557).

**Table: Applications with Implicit Support**

| Application | Implicitly Supports |
|---|---|
| 360-safeguard-update | http |
| apple-update | http |
| apt-get | http |
| as2 | http |
| avg-update | http |
| avira-antivir-update | http, ssl |
| blokus | rtmp |
| bugzilla | http |
| clubcooee | http |
| corba | http |
| cubby | http, ssl |
| dropbox | ssl |
| esignal | http |
| evernote | http, ssl |
| ezhelp | http |
| facebook | http, ssl |
| facebook-chat | jabber |
| facebook-social-plugin | http |
| fastviewer | http, ssl |
| forticlient-update | http |
| good-for-enterprise | http, ssl |
| google-cloud-print | http, ssl, jabber |
| google-desktop | http |

© Palo Alto Networks, Inc.

| Application | Implicitly Supports |
|---|---|
| google-talk | jabber |
| google-update | http |
| gotomypc-desktop-sharing | citrix-jedi |
| gotomypc-file-transfer | citrix-jedi |
| gotomypc-printing | citrix-jedi |
| hipchat | http |
| iheartradio | ssl, http, rtmp |
| infront | http |
| instagram | http, ssl |
| issuu | http, ssl |
| java-update | http |
| jepptech-updates | http |
| kerberos | rpc |
| kik | http, ssl |
| lastpass | http, ssl |
| logmein | http, ssl |
| mcafee-update | http |
| megaupload | http |
| metatrader | http |
| mocha-rdp | t_120 |
| mount | rpc |
| ms-frs | msrpc |
| ms-rdp | t_120 |
| ms-scheduler | msrpc |
| ms-service-controller | msrpc |
| nfs | rpc |
| oovoo | http, ssl |
| paloalto-updates | ssl |
| panos-global-protect | http |
| panos-web-interface | http |
| pastebin | http |
| pastebin-posting | http |

 © Palo Alto Networks, Inc.

| Application | Implicitly Supports |
|---|---|
| pinterest | http, ssl |
| portmapper | rpc |
| prezi | http, ssl |
| rdp2tcp | t_120 |
| renren-im | jabber |
| roboform | http, ssl |
| salesforce | http |
| stumbleupon | http |
| supremo | http |
| symantec-av-update | http |
| trendmicro | http |
| trillian | http, ssl |
| twitter | http |
| whatsapp | http, ssl |
| xm-radio | rtsp |

© Palo Alto Networks, Inc.

# Application Level Gateways

The Palo Alto Networks firewall does not classify traffic by port and protocol; instead it identifies the application based on its unique properties and transaction characteristics using the App-ID technology. Some applications, however, require the firewall to dynamically open *pinholes* to establish the connection, determine the parameters for the session and negotiate the ports that will be used for the transfer of data; these applications use the application-layer payload to communicate the dynamic TCP or UDP ports on which the application opens data connections. For such applications, the firewall serves as an Application Level Gateway (ALG), and it opens a pinhole for a limited time and for exclusively transferring data or control traffic. The firewall also performs a NAT rewrite of the payload when necessary.

As of Content Release version 504, the Palo Alto Networks firewall provides NAT ALG support for the following protocols: FTP, H.225, H.248, MGCP, MySQL, Oracle/SQLNet/TNS, RPC, RTSP, SCCP, SIP, and UNIStim.

When the firewall serves as an ALG for the Session Initiation Protocol (SIP), by default it performs NAT on the payload and opens dynamic pinholes for media ports. In some cases, depending on the SIP applications in use in your environment, the SIP endpoints have NAT intelligence embedded in their clients. In such cases, you might need to disable the SIP ALG functionality to prevent the firewall from modifying the signaling sessions. When SIP ALG is disabled, if App-ID determines that a session is SIP, the payload is not translated and dynamic pinholes are not opened. See Disable the SIP Application-level Gateway (ALG).

The firewall provides IPv6-to-IPv6 Network Prefix Translation (NPTv6) ALG support for the following protocols: FTP, Oracle, and RTSP. The SIP ALG is not supported for NPTv6 or NAT64.

 © Palo Alto Networks, Inc.

# Disable the SIP Application-level Gateway (ALG)

The Palo Alto Networks firewall uses the Session Initiation Protocol (SIP) application-level gateway (ALG) to open dynamic pinholes in the firewall where NAT is enabled. However, some applications—such as VoIP—have NAT intelligence embedded in the client application. In these cases, the SIP ALG on the firewall can interfere with the signaling sessions and cause the client application to stop working.

One solution to this problem is to define an Application Override Policy for SIP, but using this approach disables the App-ID and threat detection functionality. A better approach is to disable the SIP ALG, which does not disable App-ID or threat detection.

The following procedure describes how to disable the SIP ALG.

**Disable the SIP ALG**

Step 1 Select **Objects > Applications**.

Step 2 Select the **sip** application.

You can type **sip** in the **Search** box to help find the sip application.

Step 3 Select **Customize...** for **ALG** in the Options section of the Application dialog box.

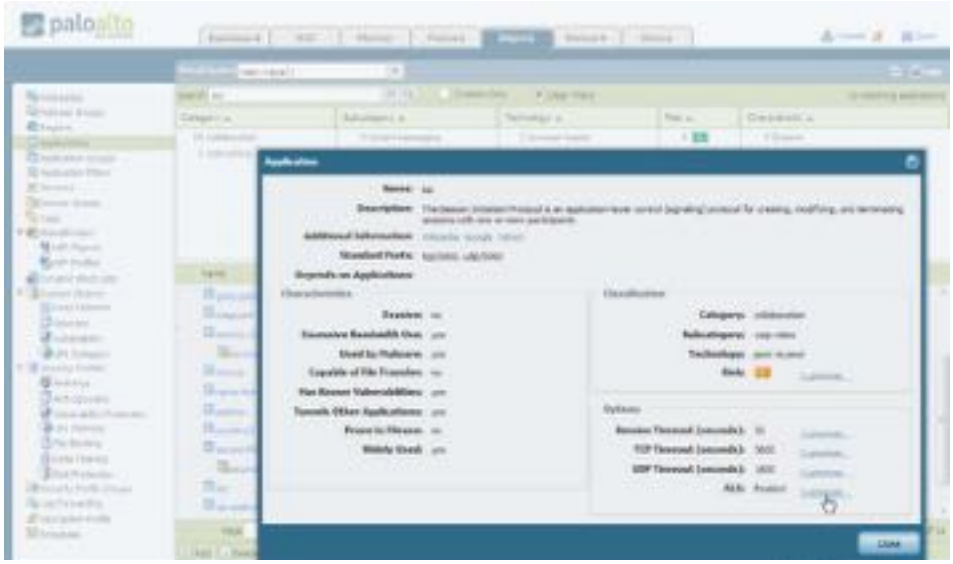

Step 4 Select the **Disable ALG** check box in the Application - sip dialog box and click **OK**.

Step 5 **Close** the Application dialog box and **Commit** the change.

© Palo Alto Networks, Inc.

# Threat Prevention

The Palo Alto Networks next-generation firewall protects and defends your network from commodity threats and advanced persistent threats (APTs). The firewall's multi-pronged detection mechanisms include a signature-based (IPS/Command and Control/Antivirus) approach, heuristics-based (bot detection) approach, sandbox-based (WildFire) approach, and Layer 7 protocol analysis-based (App-ID) approach.

Commodity threats are exploits that are less sophisticated and more easily detected and prevented using a combination of the antivirus, anti-spyware, vulnerability protection and the URL filtering/Application identification capabilities on the firewall.

Advanced threats are perpetuated by organized cyber criminals or malicious groups that use sophisticated attack vectors to target your network, most commonly for intellectual property theft and financial data theft. These threats are more evasive and require intelligent monitoring mechanisms for detailed host and network forensics on malware. The Palo Alto Networks next-generation firewall in conjunction with WildFire and Panorama provides a comprehensive solution that intercepts and break the attack chain and provides visibility to prevent security infringement on your network—including mobile and virtualized—infrastructure.


- Set Up Security Profiles and Policies
- Prevent Brute Force Attacks
- Customize the Action and Trigger Conditions for a Brute Force Signature
- Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions
- Enable Passive DNS Collection for Improved Threat Intelligence
- Use DNS Queries to Identify Infected Hosts on the Network
- Content Delivery Network Infrastructure for Dynamic Updates
- Threat Prevention Resources


© Palo Alto Networks, Inc.

# Set Up Security Profiles and Policies

The following sections provide basic threat prevention configuration examples:

- Set Up Antivirus, Anti-Spyware, and Vulnerability Protection
- Set Up Data Filtering
- Set Up File Blocking

For information on controlling web access as part of your threat prevention strategy, see URL Filtering.


© Palo Alto Networks, Inc.

# Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

The following describes the steps needed to set up the default Antivirus, Anti-Spyware, and Vulnerability Protection Security Profiles.

All anti-spyware and vulnerability protection signatures have a default action defined by Palo Alto Networks. You can view the default action by navigating to **Objects > Security Profiles > Anti-Spyware** or **Objects > Security Profiles >Vulnerability Protection** and then selecting a profile. Click the **Exceptions** tab and then click **Show all signatures** and you will see a list of the signatures with the default action in the Action column. To change the default action, you must create a new profile and then create rules with a non-default action, and/or add individual signature exceptions to **Exceptions** in the profile.

| **Set up Antivirus/Anti-Spyware/Vulnerability Protection** | | |
|---|---|---|
| Step 1 | Verify that you have a Threat Prevention license. | • The Threat Prevention license bundles the antivirus, anti-spyware, and the vulnerability protection features in one license.<br>• Select **Device > Licenses** to verify that the **Threat Prevention** license is installed and check the expiration date. |
| Step 2 | Download the latest antivirus threat signatures. | 1. Select **Device > Dynamic Updates** and click **Check Now** at the bottom of the page to retrieve the latest signatures.<br>In the **Actions** column, click **Download** to install the latest Antivirus and Applications and Threats signatures. |
| Step 3 | Schedule signature updates. | 1. From **Device > Dynamic Updates**, click the text to the right of **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.<br>2. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select **Download Only**, you would need to manually go in and click the **Install** link in the **Action** column to install the signature. When you click **OK**, the update is scheduled. No commit is required.<br>3. (Optional) You can also enter the number of hours in the **Threshold** field to indicate the minimum age of a signature before a download will occur. For example, if you entered **10**, the signature must be at least 10 hours old before it will be downloaded, regardless of the schedule.<br>4. In an HA configuration, you can also click the **Sync To Peer** option to synchronize the content update with the HA peer after download/install. This will not push the schedule settings to the peer device, you need to configure the schedule on each device. |

© Palo Alto Networks, Inc.

**Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued)**

### Best Practices for Antivirus Schedules

The general recommendation for antivirus signature update schedules is to perform a **download-and-install** on a daily basis for antivirus and weekly for applications and vulnerabilities.

### Recommendations for HA Configurations:

- **Active/Passive HA**—If the MGT port is used for antivirus signature downloads, you should configure a schedule on both devices and both devices will download/install independently. If you are using a data port for downloads, the passive device will not perform downloads while it is in the passive state. In this case you would set a schedule on both devices and then select the **Sync To Peer** option. This will ensure that whichever device is active, the updates will occur and will then push to the passive device.
- **Active/Active HA**—If the MGT port is used for antivirus signature downloads on both devices, then schedule the download/install on both devices, but do not select the **Sync To Peer** option. If you are using a data port, schedule the signature downloads on both devices and select **Sync To Peer**. This will ensure that if one device in the active/active configuration goes into the active-secondary state, the active device will download/install the signature and will then push it to the active-secondary device.

| | |
|---|---|
| Step 4 Attach the security profiles to a security policy. | 1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab.<br>2. In **Profile Settings**, click the drop-down next to each security profile you would like to enable. In this example we choose default for **Antivirus, Vulnerability Protection,** and **Anti-Spyware**.<br>If no security profiles have been previously defined, select **Profiles** from the **Profile Type** drop-down. You will then see the list of options to select the security profiles.<br> 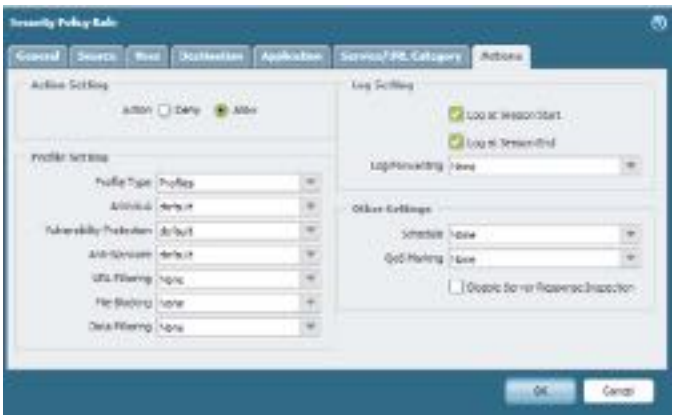 |
| Step 5 Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

## Set Up Data Filtering

The following describes the steps needed to configure a data filtering profile that will detect Social Security Numbers and a custom pattern identified in .doc and .docx documents.

| Data Filtering Configuration Example | | |
|---|---|---|
| Step 1 | Create a Data Filtering security profile. | 1. Select **Objects > Security Profiles > Data Filtering** and click **Add**.<br>2. Enter a **Name** and a **Description** for the profile. In this example the name is *DF_Profile1* with the description *Detect Social Security Numbers*.<br>3. (Optional) If you want to collect data that is blocked by the filter, select the **Data Capture** check box.<br> You must set a password as described in Step 2 if you are using the data capture feature. |
| Step 2 | (Optional) Secure access to the data filtering logs to prevent other administrators from viewing sensitive data.<br>When you enable this option, you will be prompted for the password when you view logs in **Monitor > Logs > Data Filtering**. | 1. Select **Device > Setup > Content-ID**.<br>2. Click **Manage Data Protection** in the Content-ID Features section.<br>3. Set the password that will be required to view the data filtering logs. |

**Data Filtering Configuration Example (Continued)**

Step 3 Define the data pattern that will be used in the Data Filtering Profile.

In this example, we will use the keyword **confidential** and will set the option to search for SSN numbers with dashes (Example - 987-654-4320).

It is helpful to set the appropriate thresholds and define keywords within documents to reduce false positives.

1. From the Data Filtering Profile page click **Add** and select **New** from the **Data Pattern** drop-down. You can also configure data patterns from **Objects > Custom Signatures > Data Patterns**.
2. For this example, name the Data Pattern signature **Detect SS Numbers** and add the description **Data Pattern to detect Social Security numbers**.
3. In the **Weight** section for **SSN#** enter **3**. See Weight and Threshold Values for more details.

4. (Optional) You can also set **Custom Patterns** that will be subject to this profile. In this case, you specify a pattern in the custom patterns **Regex** field and set a weight. You can add multiple match expressions to the same data pattern profile. In this example, we will create a **Custom Pattern** named **SSN_Custom** with a custom pattern of **confidential** (the pattern is case sensitive) and use a weight of **20**. The reason we use the term **confidential** in this example is because we know that our social security Word docs contain this term, so we define that specifically.

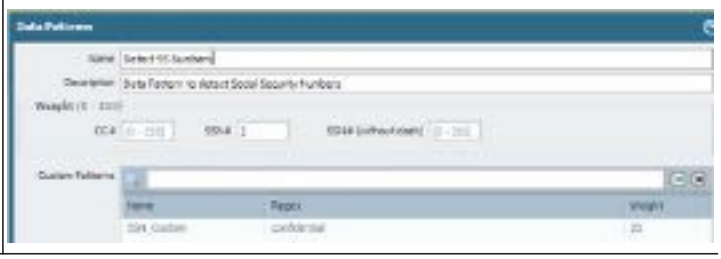

 © Palo Alto Networks, Inc.

| Data Filtering Configuration Example (Continued) | | |
|---|---|---|
| Step 4 | Specify which applications to filter and set the file types. | 1. Set **Applications** to **Any.** This will detect any supported application such as: web-browsing, FTP, or SMTP. If you want to narrow down the application, you can select it from the list. For applications such as Microsoft Outlook Web App that uses SSL, you will need to enable decryption. Also make sure you understand the naming for each application. For example, Outlook Web App, which is the Microsoft name for this application is identified as the application **outlook-web** in the PAN-OS list of applications. You can check the logs for a given application to identify the name defined in PAN-OS.<br>2. Set **File Types** to **doc** and **docx** to only scan doc and docx files. |
| Step 5 | Specify the direction of traffic to filter and the threshold values. | 1. Set the **Direction** to **Both**. Files that are uploaded or downloaded will be scanned.<br>2. Set the **Alert Threshold** to **35.** In this case, an alert will be triggered if 5 instances of Social Security Numbers exist and 1 instance of the term **confidential** exists. The formula is 5 SSN instances with a weight of 3 = 15 plus 1 instance of the term confidential with a weight of 20 = 35.<br>3. Set the **Block Threshold** to **50**. The file will be blocked if the threshold of 50 instances of a SSN and/or the term confidential exists in the file. In this case, if the doc contained 1 instance of the word **confidential** with a weight of 20 that equals 20 toward the threshold, and the doc has 15 Social Security Numbers with a weight of 3 that equals 45. Add 20 and 45 and you have 65, which will exceed the block threshold of 50.<br> 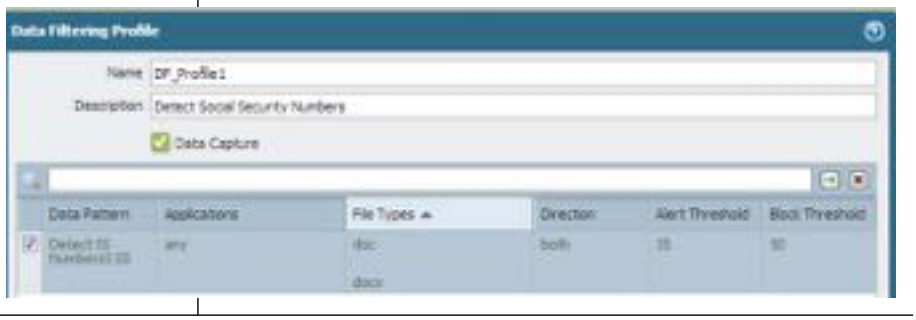 |

**Data Filtering Configuration Example (Continued)**

| | | |
|---|---|---|
| Step 6 | Attach the Data Filtering profile to the security rule. | 1. Select **Policies > Security** and select the security policy rule to which to apply the profile.<br>2. Click the security policy rule to modify it and then click the **Actions** tab. In the **Data Filtering** drop-down, select the new data filtering profile you created and then click **OK** to save. In this example, the data filtering rule name is **DF_Profile1**.<br> 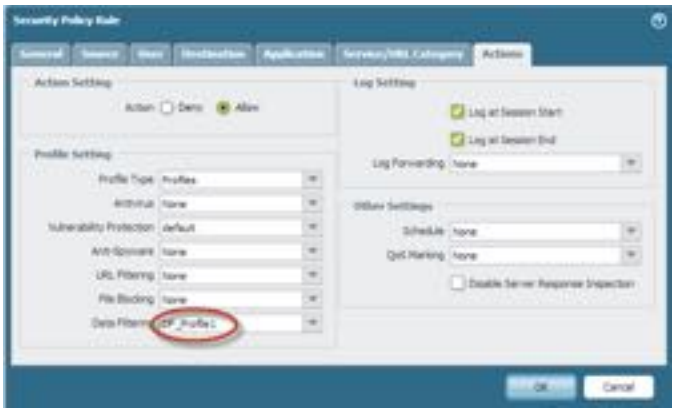 |
| Step 7 | **Commit** the configuration. | |
| Step 8 | Test the data filtering configuration.<br><br>If you have problems getting Data Filtering to work, you can check the Data Filtering log or the Traffic log to verify the application that you are testing with and make sure your test document has the appropriate number of unique Social Security Number instances. For example, an application such as Microsoft Outlook Web App may seem to be identified as web-browsing, but if you look at the logs, the application is **outlook-web**. Also increase the number of SSNs, or your custom pattern to make sure you are hitting the thresholds. | When testing, you must use real Social Security Numbers and each number must be unique. Also, when defining Custom Patterns as we did in this example with the word **confidential**, the pattern is case sensitive. To keep your test simple, you may want to just test using a data pattern first, then test the SSNs.<br>1. Access a client PC in the trust zone of the firewall and send an HTTP request to upload a .doc or .docx file that contains the exact information you defined for filtering.<br>2. Create a Microsoft Word document with one instance of the term confidential and five Social Security numbers with dashes.<br>3. Upload the file to a website. Use an HTTP site unless you have decryption configured, in which case you can use HTTPS.<br>4. Select **Monitoring > Logs > Data Filtering** logs.<br>5. Locate the log that corresponds to the file you just uploaded. To help filter the logs, use the source of your client PC and the destination of the web server. The action column in the log will show **reset-both**. You can now increase the number of Social Security Numbers in the document to test the block threshold. |

 © Palo Alto Networks, Inc.

## Set Up File Blocking

This example will describe the basic steps needed to set up file blocking. In this configuration, we will configure the options needed to prompt users to continue before downloading .exe files from websites. When testing this example, be aware that you may have other systems between you and the source that may be blocking content.

| **Configure File Blocking** | | |
|---|---|---|
| Step 1 | Create the file blocking profile. | 1. Select **Objects > Security Profiles > File Blocking** and click **Add**.<br>2. Enter a **Name** for the file blocking profile, for example *Block_EXE*. Optionally enter a **Description**, such as *Block users from downloading exe files from websites.* |
| Step 2 | Configure the file blocking options. | 1. Click **Add** to define the profile settings.<br>2. Enter a **Name**, such as *BlockEXE*.<br>3. Set the **Applications** for filtering, for example **web-browsing**.<br>4. Set **File Types** to **exe**.<br>5. Set the **Direction** to **download**.<br>6. Set the **Action** to **continue**. By choosing the continue option, users will be prompted with a response page prompting them to click continue before the file will be downloaded.<br> <br>7. Click **OK** to save the profile. |
| Step 3 | Apply the file blocking profile to a security policy. | 1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Set Up Basic Security Policies.<br>2. Click the **Actions** tab within the policy rule.<br>3. In the Profile Settings section, click the drop-down and select the file blocking profile you configured. In this case, the profile name is *Block_EXE*.<br>4. **Commit** the configuration.<br>If no security profiles have been previously defined, select the **Profile Type** drop-down and select **Profiles**. You will then see the list of options to select the security profiles. |

**Configure File Blocking (Continued)**

Step 4 To test your file blocking configuration, access a client PC in the trust zone of the firewall and attempt to download an .exe file from a website in the untrust zone. A response page should display. Click **Continue** to download the file. You can also set other actions, such as alert or block, which will not provide a continue page to the user. The following shows the default response page for File Blocking:

**Example: Default File Blocking Response Page**

Step 5 (Optional) Define custom file blocking response pages (**Device > Response Pages**). This allows you to provide more information to users when they see a response page. You can include information such as company policy information and contact information for a Helpdesk.

When you create a file blocking profile with the action continue, you can only choose the application web-browsing. If you choose any other application, traffic that matches the security policy will not flow through the firewall due to the fact that the users will not be prompted with a continue page. Also, if the website uses HTTPS, you will need to have a decryption policy in place.

You may want to check your logs to confirm what application is being used when testing this feature. For example, if you are using Microsoft SharePoint to download files, even though you are using a web-browser to access the site, the application is actually `sharepoint-base`, or `sharepoint-document`. You may want to set the application type to **Any** for testing.

 © Palo Alto Networks, Inc.

# Prevent Brute Force Attacks

A brute force attack uses a large volume of requests/responses from the same source or destination IP address to break into a system. The attacker employs a trial-and-error method to guess the response to a challenge or a request.

The Vulnerability Protection profile on the firewall includes signatures to protect you from brute force attacks. Each signature has an ID, Threat Name, Severity and is triggered when a pattern is recorded. The pattern specifies the conditions and interval at which the traffic is identified as a brute-force attack; some signatures are associated with another child signature that is of a lower severity and specifies the pattern to match against. When a pattern matches against the signature or child signature, it triggers the default action for the signature.

To enforce protection:

- Attach the vulnerability profile to a security rule. See Set Up Antivirus, Anti-Spyware, and Vulnerability Protection.
- Install content updates that include new signatures to protect against emerging threats. See Manage Content Updates.

# Customize the Action and Trigger Conditions for a Brute Force Signature

The firewall includes two types of predefined brute force signatures—parent signature and child signature. A child signature is a single occurrence of a traffic pattern that matches the signature. A parent signature is associated with a child signature and is triggered when multiple events occur within a time interval and match the traffic pattern defined in the child signature.

Typically, a child signature is of default action *allow* because a single event is not indicative of an attack. In most cases, the action for a child signature is set to allow so that legitimate traffic is not blocked and threat logs are not generated for non-noteworthy events. Therefore, Palo Alto Networks recommends that you only change the default action after careful consideration.

In most cases, the brute force signature is a noteworthy event because of its recurrent pattern. If you would like to customize the action for a brute-force signature, you can do one of the following:

- Create a rule to modify the default action for all signatures in the brute force category. You can define the action to allow, alert, block, reset, or drop the traffic.
- Define an exception for a specific signature. For example, you can search for a CVE and define an exception for it.

  For a parent signature, you can modify both the trigger conditions and the action; for a child signature only the action can be modified.

To effectively mitigate an attack, the **block-ip address** action is recommended over the drop or reset action for most brute force signatures.

| **Customize the Threshold and Action for a Signature** | | |
|---|---|---|
| Step 1 | Create a new Vulnerability Protection profile. | 1. Select **Objects > Security Profiles > Vulnerability Protection**.<br>2. Click **Add** and enter a **Name** for the Vulnerability Protection profile. |
| Step 2 | Create a rule that defines the action for all signatures in a category.<br> | 1. Select **Rules**, click **Add** and enter a **Name** for the rule.<br>2. Set the **Action**. In this example, it is set to **Block**.<br>3. Set **Category** to **brute-force**.<br>4. (Optional) If blocking, specify whether to block server or client, the default is any.<br>5. See Step 3 to customize the action for a specific signature.<br>6. See Step 4 to customize the trigger threshold for a parent signature.<br>7. Click **OK** to save the rule and the profile. |


© Palo Alto Networks, Inc.

**Customize the Threshold and Action for a Signature**

Step 3 (Optional) Customize the action for a specific signature.

1. Select **Exceptions** and click **Show all signatures** to find the signature you want to modify.

   To view all the signatures in the brute-force category, search for (category contains 'brute-force').
2. To edit a specific signature, click the predefined default action in the **Action** column.

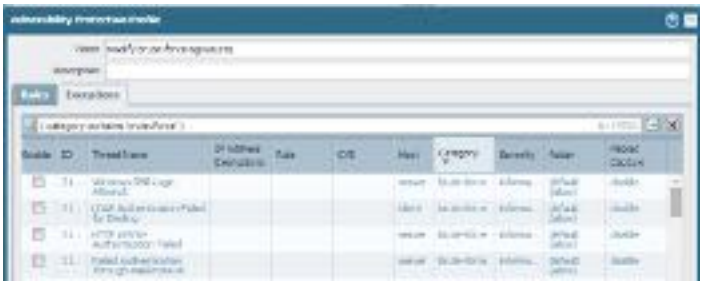

3. Set the action to **allow**, **alert** or **block-ip**.

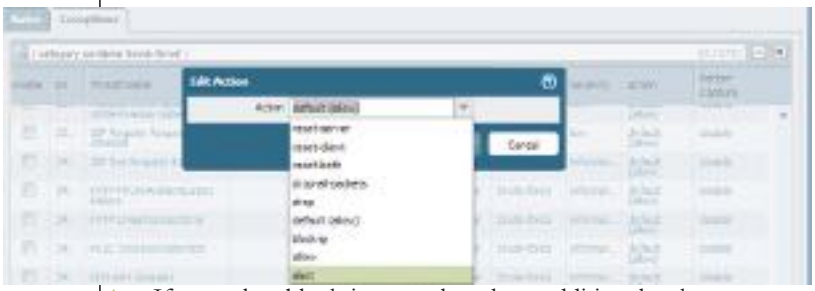

4. If you select block-ip, complete these additional tasks:
   a. Specify the **Time** period (in seconds) after which to trigger the action.
   b. In the **Track By** field, define whether to block the IP address by **IP source** or by **IP source and destination**.

5. Click **OK**
6. For each modified signature, select the check box in the **Enable** column.
7. Click **OK**.

| Customize the Threshold and Action for a Signature | | |
|---|---|---|
| Step 4 | Customize the trigger conditions for a parent signature.<br><br>A parent signature that can be edited is marked with this icon: .<br><br>In this example, the search criteria was brute force category and CVE-2008-1447. | 1. Click to edit the time attribute and the aggregation criteria for the signature.<br> <br>2. To modify the trigger threshold specify the **Number of Hits** *per x* **seconds**.<br>3. Specify whether to aggregate the number of hits by **source**, **destination** or by **source and destination**.<br>4. Click **OK**. |
| Step 5 | Attach this new profile to a security rule. | |
| Step 6 | Save your changes. | 1. Click **Commit**. |

 © Palo Alto Networks, Inc.

# Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions

To monitor and protect your network from most Layer 4 and Layer 7 attacks, here are a few recommendations.

- Upgrade to the most current PAN-OS software version and content release version to ensure that you have the latest security updates. See Manage Content Updates and Install Software Updates.
- For servers, create security policy rules to only allow the application(s) that you sanction on each server. Verify that the standard port for the application matches the listening port on the server. For example, to ensure that only SMTP traffic is allowed to your email server set the Application to smtp and the Service to application-default. If your server uses only a subset of the standard ports (for example, if your SMTP server uses only port 587 while the smtp application has standard ports defined as 25 and 587), you should create a new custom service that only includes port 587 and use that new service in your security policy rule instead of using application-default. Additionally, make sure to restrict access to specific source and destinations zones and sets of IP addresses.
- Attach the following security profiles to your security policies to provide signature-based protection.
  - Create a Vulnerability Protection profile to block all vulnerabilities with severity low and higher.
  - Create an Anti-Spyware profile to block all spyware.
  - Create an antivirus profile to block all content that matches an antivirus signature.
- Block all unknown applications/traffic using security policy. Typically, the only applications that are classified as unknown traffic are internal or custom applications on your network, or potential threats. Because unknown traffic can be a non-compliant application or protocol that is anomalous or abnormal, or a known application that is using non-standard ports, unknown traffic should be blocked. See Manage Custom or Unknown Applications.
- Create a file blocking profile that blocks Portable Executable (PE) file types for Internet-based SMB (Server Message Block) traffic from traversing the trust to untrust zones, (ms-ds-smb applications).
- Create a zone protection profile that is configured to protect against packet-based attacks:
  - Remove TCP timestamps on SYN packets before the firewall forwards the packet—When you remove the TCP timestamp option in a SYN packet, the TCP stack on both ends of the TCP connection will not support TCP timestamps. Therefore, by disabling the TCP timestamp for a SYN packet, you can prevent an attack that uses different timestamps on multiple packets for the same sequence number.
  - Drop malformed packets.
  - Drop mismatched and overlapping TCP segments—By deliberately constructing connections with overlapping but different data in them, attackers can attempt to cause misinterpretation of the intent of the connection. This can be used to deliberately induce false positives or false negatives. An attacker can use IP spoofing and sequence number prediction to intercept a user's connection and inject his/her own data into the connection. PAN-OS uses this field to discard such frames with mismatched and overlapping data. The scenarios where the received segment will be discarded are when the segment received is contained within another segment, the segment received overlaps with part of another segment, or the segment completely contains another segment.

© Palo Alto Networks, Inc.

- Verify that support for IPv6 is enabled, if you have configured IPv6 addresses on your network hosts. (**Network > Interfaces > Ethernet> IPv6**)

  This allows access to IPv6 hosts and filters IPv6 packets that are encapsulated in IPv4 packets. Enabling support for IPv6 prevents IPv6 over IPv4 multicast addresses from being leveraged for network reconnaissance.

- Enable support for multicast traffic so that the firewall can enforce policy on multicast traffic. (**Network > Virtual Router > Multicast**)

- Enable the following CLI command to clear the URG bit flag in the TCP header and disallow out-of-band processing of packets.

  The urgent pointer in the TCP header is used to promote a packet for immediate processing by removing it from the processing queue and expediting it through the TCP/IP stack on the host. This process is called out-of-band processing. Because the implementation of the urgent pointer varies by host, to eliminate ambiguity, use the following CLI command to disallow out-of-band processing; the out-of-band byte in the payload becomes part of the payload and the packet is not processed urgently. Making this change allows you to remove ambiguity in how the packet is processed on the firewall and the host, and the firewall sees the exact same stream in the protocol stack as the host for whom the packet is destined.

  ```
  set deviceconfig setting tcp urgent-data clear
  ```

- If you configure the firewall to clear the URG bit flag and the packet has no other flags set in the TCP header, use the following CLI command to configure the firewall to drop packets with no flags:

  ```
  set deviceconfig setting tcp drop-zero-flag yes
  ```

- Enable the following CLI command for disabling the bypass-exceed-queue.

  The bypass exceed queue is required for out of order packets. This scenario is most common in an asymmetric environment where the firewall receives packets out of order. For identification of certain applications (App-ID) the firewall performs heuristic analysis. If the packets are received out of order, the data must be copied to a queue in order to complete the analysis for the application.

  ```
  set deviceconfig setting application bypass-exceed-queue no
  ```

- Enable the following CLI commands for disabling the inspection of packets when the out-of-order packet limit is reached. The Palo Alto Networks firewall can collect up to 64 out-of-order packets per session. This counter identifies that packets have exceeded the 64-packet limit. When the bypass setting is set to **no**, the device drops the out-of-order packets that exceed the 64-packet limit. A commit is required.

  ```
  set deviceconfig setting tcp bypass-exceed-oo-queue no
  set deviceconfig setting ctd tcp-bypass-exceed-queue no
  set deviceconfig setting ctd udp-bypass-exceed-queue no
  ```

- Enable the following CLI commands for checking the TCP timestamp. The TCP timestamp records when the segment was sent and allows the firewall to verify that the timestamp is valid for that session. Packets with invalid timestamps are dropped with this setting is enabled.

  ```
  set deviceconfig setting tcp check-timestamp-option yes
  ```

 © Palo Alto Networks, Inc.

- Disable the HTTP Range option. The HTTP Range option allows a client to fetch part of a file only. When a next-generation firewall in the path of a transfer identifies and drops a malicious file, it terminates the TCP session with a RST packet. If the web browser implements the HTTP Range option, it can start a new session to fetch only the remaining part of the file. This prevents the firewall from triggering the same signature again due to the lack of context into the initial session, while at the same time allowing the web browser to reassemble the file and deliver the malicious content. To prevent this, disable the HTTP Range option as follows:

```
set deviceconfig setting ctd skip-block-http-range no
```

# Enable Passive DNS Collection for Improved Threat Intelligence

Passive DNS is an opt-in feature that enables the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis in order to improve threat intelligence and threat prevention capabilities. The data collected includes non-recursive (i.e. originating from the local recursive resolver, not individual clients) DNS query and response packet payloads. Data submitted via the Passive DNS Monitoring feature consists solely of mappings of domain names to IP addresses. Palo Alto Networks retains no record of the source of this data and does not have the ability to associate it with the submitter at a future date.

The Palo Alto Networks threat research team uses this information to gain insight into malware propagation and evasion techniques that abuse the DNS system. Information gathered through this data collection is used to improve accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control signatures, and WildFire.

DNS responses are only forwarded to the Palo Alto Networks and will only occur when the following requirements are met:

- DNS response bit is set
- DNS truncated bit is not set
- DNS recursive bit is not set
- DNS response code is 0 or 3 (NX)
- DNS question count bigger than 0
- DNS Answer RR count is bigger than 0 or if it is 0, the flags need to be 3 (NX)
- DNS query record type are A, NS, CNAME, AAAA, MX

Passive DNS monitoring is disabled by default, but it is recommended that you enable it to facilitate enhanced threat intelligence. Use the following procedure to enable Passive DNS:

| **Enable Passive DNS** | |
|---|---|
| Step 1 | Select **Objects > Security Profiles > Anti-Spyware**. |
| Step 2 | Select an existing profile to modify it or configure a new profile.<br> The Anti-Spyware profile must be attached to a security policy that governs your DNS server's external DNS traffic. |
| Step 3 | Select the **DNS Signatures** tab and click the **Enable Passive DNS Monitoring** check box. |
| Step 4 | Click **OK** and then **Commit**. |

 © Palo Alto Networks, Inc.

# Use DNS Queries to Identify Infected Hosts on the Network

The DNS sinkhole action in Anti-Spyware profiles enables the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. This allows you to identify hosts on your network that have been infected with malware. The following topics describe the DNS sinkhole action and provide instructions for enabling it and monitoring logs to identify infected hosts.

- DNS Sinkholing
- Configure DNS Sinkholing
- Identify Infected Hosts

## DNS Sinkholing

DNS sinkholing helps you to identify infected hosts on the protected network using DNS traffic in situations where the firewall cannot see the infected client's DNS query (that is, the firewall cannot see the originator of the DNS query). In a typical deployment where the firewall is north of the local DNS server, the threat log will identify the local DNS resolver as the source of the traffic rather than the actual infected host. Sinkholing malware DNS queries solves this visibility problem by forging responses to the client host queries directed at malicious domains, so that clients attempting to connect to malicious domains (for command-and-control, for example) will instead attempt to connect to a sinkhole IP address you define as illustrated in Configure DNS Sinkholing. Infected hosts can then be easily identified in the traffic logs because any host that attempts to connect to the sinkhole IP address are most likely infected with malware.

**Figure: DNS Sinkholing Example**

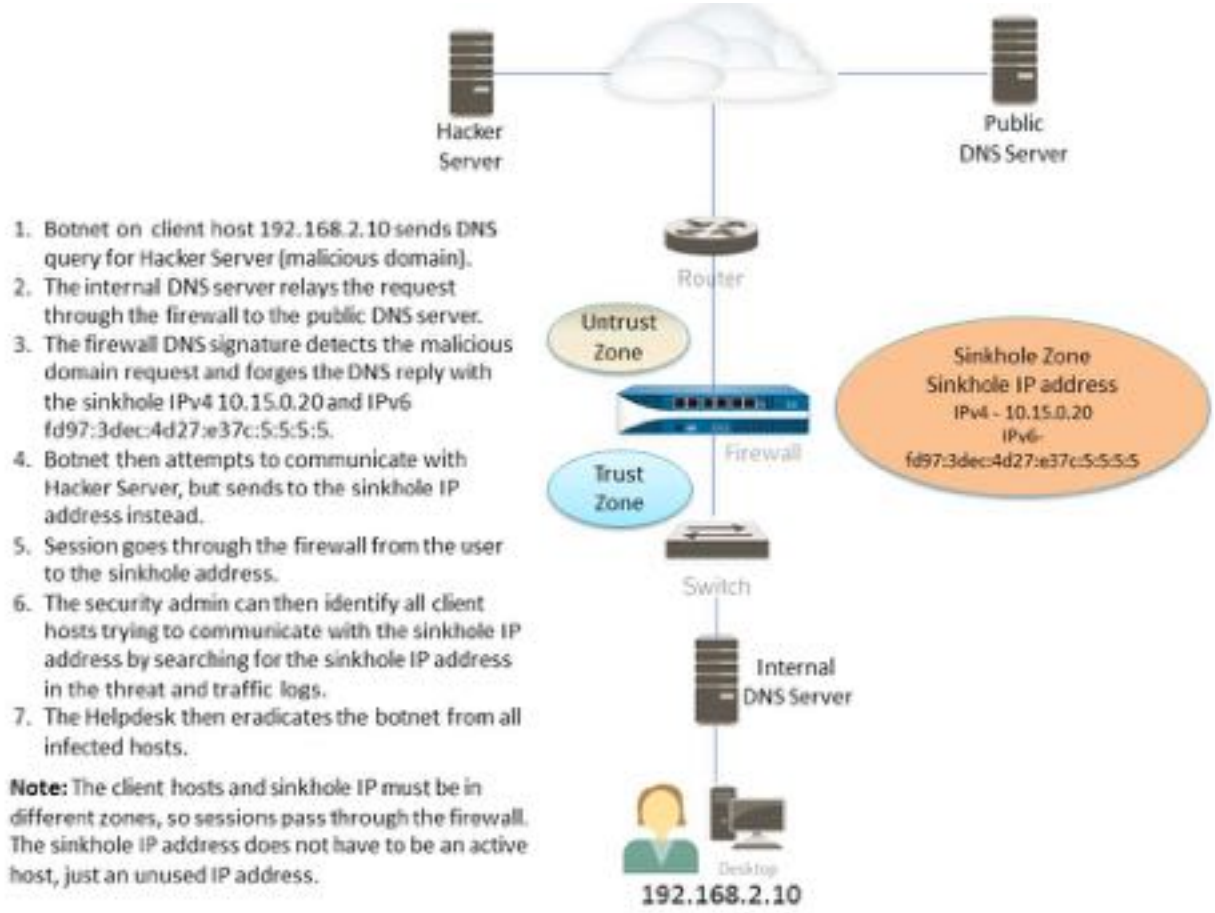


 © Palo Alto Networks, Inc.

## Configure DNS Sinkholing

To enable DNS Sinkholing, you must enable the action in an Anti-Spyware profile and attach the profile to a security rule. When a client host attempts to access a malicious domain, the firewall forges the destination IP address in the packet using the IP address you configure as the DNS sinkhole address.

| **Configure DNS Sinkholing** | | |
|---|---|---|
| Step 1 | Obtain both an IPv4 and IPv6 address to use as the sinkhole IP addresses.<br><br>The DNS sinkhole address must be in a different zone than the client hosts to ensure that when an infected host attempts to start a session with the sinkhole IP address, it will be routed through the firewall. The reason both IPv4 and IPv6 are needed is because malicious software may perform DNS queries using one or both of these protocols.<br><br> This sinkhole addresses must be reserved for this purpose and do not need to be assigned to a physical host. You can optionally use a honey-pot server as a physical host to further analyze the malicious traffic. | The configuration steps that follow use the following example DNS sinkhole addresses:<br><br>IPv4 DNS sinkhole address—10.15.0.20<br><br>IPv6 DNS sinkhole address—fd97:3dec:4d27:e37c:5:5:5:5 |
| Step 2 | Configure the sinkhole interface and zone.<br><br>Traffic from the zone where the client hosts reside must route to the zone where the sinkhole IP address is defined, so traffic will be logged.<br><br> Use a dedicated zone for sinkhole traffic, because the infected host will be sending traffic to this zone. | 1. Select **Network > Interfaces** and select an interface to configure as your sinkhole interface.<br>2. In the **Interface Type** drop-down, select **Layer3**.<br>3. To add an IPv4 address, select the **IPv4** tab and select **Static** and then click **Add**. In this example, add 10.15.0.20 as the IPv4 DNS sinkhole address.<br>4. Select the **IPv6** tab and click **Static** and then click **Add** and enter an IPv6 address and subnet mask. In this example, enter fd97:3dec:4d27:e37c::/64 as the IPv6 sinkhole address.<br>5. Click **OK** to save.<br>6. To add a zone for the sinkhole, select **Network > Zones** and click **Add**.<br>7. Enter zone **Name**.<br>8. In the **Type** drop-down select **Layer3**.<br>9. In the **Interfaces** section, click **Add** and add the interface you just configured.<br>10. Click **OK**. |

| Configure DNS Sinkholing (Continued) | | |
|---|---|---|
| Step 3 | Enable DNS sinkholing on the Anti-Spyware profile. | 1. Select **Objects > Security Profiles > Anti-Spyware**.<br>2. Modify an existing profile, or select one of the existing defaults and clone it.<br>3. **Name** the profile and then select the **DNS Signatures** tab.<br>4. In the **Action on DNS queries** drop-down, select **sinkhole**.<br>5. In the **Sinkhole IPv4** field enter the sinkhole IPv4 sinkhole address you configured in Step 2 (10.15.0.20 in this example).<br>6. In the **Sinkhole IPv6** field enter the IPv6 sinkhole address you configured in Step 2 (fd97:3dec:4d27:e37c:5:5:5:5 in this example).<br>The default sinkhole address is the loopback address (127.0.0.1 for IPv4 and ::1 for IPv6).<br>7. (Optional) In the **Packet Capture** drop-down, select **single-packet** or **extended-capture**. The single-packet option will capture the first packet of the session or select extended to set between 1-50 packets. You can then use the packet captures for further analysis.<br>8. Click **OK** to save the profile. |
| Step 4 | Edit the security policy rule that allows traffic from client hosts in the trust zone to the untrust zone to include the sinkhole zone as a destination and attach the Anti-Spyware profile.<br>To ensure that you are identifying traffic from infected hosts, make these changes to the security rule(s) that allow traffic from client hosts in the trust zone to the untrust zone. By adding the sinkhole zone as a destination on the rule, you enable infected clients to send bogus DNS queries to the DNS sinkhole. | 1. Select **Policies > Security**.<br>2. Select an existing rule that allows traffic from the client host zone to the untrust zone.<br>3. On the **Destination** tab, **Add** the Sinkhole zone. This allows client host traffic to flow to the sinkhole zone.<br>4. On the **Actions** tab, select the **Log at Session Start** check box to enable logging. This will ensure that traffic from client hosts in the Trust zone will be logged when accessing the Untrust or Sinkhole zones.<br>5. In the **Profile Setting** section, select the **Anti-Spyware** profile in which you enabled DNS sinkholing.<br>6. Click **OK** to save the security rule and then **Commit**. |

 © Palo Alto Networks, Inc.

| Configure DNS Sinkholing (Continued) | | |
|---|---|---|
| Step 5 | To ensure that you will be able to identify infected hosts, verify that traffic going from the client host in the Trust zone to the new Sinkhole zone is being logged.<br><br>In this example, the infected client host is 192.168.2.10 and the Sinkhole IPv4 address is 10.15.0.20. | 1. From a client host in the trust zone, open a command prompt and run the following command:<br>`C:\>ping <sinkhole address>`<br>The following example output shows the ping request to the DNS sinkhole address at 10.15.0.2 and the result, which is `Request timed out` because in this example the sinkhole IP address is not assigned to a physical host:<br>`C:\>ping 10.15.0.20`<br>`Pinging 10.15.0.20 with 32 bytes of data:`<br>`Request timed out.`<br>`Request timed out.`<br>`Ping statistics for 10.15.0.20:`<br>`Packets: Sent = 4, Received = 0, Lost = 4 (100% loss)`<br>2. On the firewall, select **Monitor > Logs > Traffic** and find the log entry with the Source 192.168.2.10 and Destination 10.15.0.20. This will confirm that the traffic to the sinkhole IP address is traversing the firewall zones.<br> You can search and/or filter the logs and only show logs with the destination 10.15.0.20. To do this, click the IP address (10.15.0.20) in the **Destination** column, which will add the filter (addr.dst in 10.15.0.20) to the search field. Click the Apply Filter icon to the right of the search field to apply the filter. |
| Step 6 | Identify a malicious domain that you can use to verify that the DNS sinkhole functionality is configured properly.<br><br>You must test this feature using a malicious domain that is included in the firewall's current antivirus signature database. The DNS Signatures used to identify malicious domains is only part of the full antivirus signature database, which contains hundreds of thousands of signatures. | To find a malicious domain for testing:<br>1. Select **Device > Dynamic Updates** and in the **Antivirus** section click the **Release Notes** link for the current antivirus DB that is installed. You can also find the antivirus release notes on the support site in Dynamic Updates. In most cases, the signature update is an incremental update, so only new viruses and DNS signatures are listed. There are many antivirus signatures and DNS signatures that will already be installed on the firewall.<br>2. In the second column of the release note, locate a line item with a domain extension (for example, .com, .edu, or .net). The left column displays the domain name. For example, in Antivirus release 1117-1560, there is an item in the left column named **tbsbana** and the right column lists **net**.<br>The following shows the content in the release note for this line item:<br>`conficker:tbsbana1 variants: net`<br>Because this domain shows up in the current database, it will work for testing. |

| Configure DNS Sinkholing (Continued) | |
|---|---|
| Step 7 Test the sinkhole action<br><br>This is similar to the action that would be performed if the client host was infected and the malicious application was attempting to reach a hacker server using DNS queries. | 1. From the client host, open a command prompt.<br>2. Perform an NSLOOKUP to a URL that you identified as a known malicious domain in Step 6.<br><br>For example, using the URL `track.bidtrk.com`:<br>`C:\>nslookup track.bidtrk.com`<br>`Server: my-local-dns.local`<br>`Address: 10.0.0.222`<br>`Non-authoritative answer:`<br>`Name: track.bidtrk.com.org`<br>`Addresses: fd97:3dec:4d27:e37c:5:5:5:5`<br>`10.15.0.20`<br><br>In the output, note that the NSLOOKUP to the malicious domain has been forged using the sinkhole IP addresses that we configured (10.15.0.20). Because the domain matched a malicious DNS signature, the sinkhole action was performed.<br>3. Select **Monitor > Logs > Threat** and locate the corresponding threat log entry to verify that the correct action was taken on the NSLOOKUP request.<br>4. Perform a ping to `track.bidtrk.com`, which will generate network traffic to the sinkhole address. |

 © Palo Alto Networks, Inc.

## Identify Infected Hosts

After you have configured DNS sinkholing and verified that traffic to a malicious domain goes to the sinkhole address, you should regularly monitor traffic to the sinkhole address, so that you can track down the infected hosts and eliminate the threat.

| DNS Sinkhole Verification and Reporting | | |
|---|---|---|
| Step 1 | Use App Scope to identify infected client hosts. | 1. Select **Monitor > App Scope** and select **Threat Monitor**.<br>2. Click the **Show spyware** button along the top of the display page.<br>3. Select a time range.<br>The following screenshot shows three instances of Suspicious DNS queries, which were generated when the test client host performed an NSLOOKUP on a known malicious domain. Click the graph to see more details about the event.<br> |

**DNS Sinkhole Verification and Reporting (Continued)**

Step 2 Configure a custom report to identify all client hosts that have sent traffic to the sinkhole IP address, which is 10.15.0.20 in this example.

Forward to an SNMP manager, Syslog server and/or Panorama to enable alerts on these events.

In this example, the infected client host performed an NSLOOKUP to a known malicious domain that is listed in the Palo Alto Networks DNS Signature database. When this occurred, the query was sent to the local DNS server, which then forwarded the request through the firewall to an external DNS server. The firewall security policy with the Anti-Spyware profile configured matched the query to the DNS Signature database, which then forged the reply using the sinkhole address of 10.15.0.20 and fd97:3dec:4d27:e37c:5:5:5:5. The client attempts to start a session and the traffic log records the activity with the source host and the destination address, which is now directed to the forged sinkhole address.

Viewing the traffic log on the firewall allows you to identify any client host that is sending traffic to the sinkhole address. In this example, the logs show that the source address 192.168.2.10 sent the malicious DNS query. The host can then be found and cleaned. Without the DNS sinkhole option, the administrator would only see the local DNS server as the system that performed the query and would not see the client host that is infected. If you attempted to run a report on the threat log using the action "Sinkhole", the log would show the local DNS server, not the infected host.

1. Select **Monitor > Manage Custom Reports**.
2. Click **Add** and **Name** the report.
3. Define a custom report that captures traffic to the sinkhole address as follows:
   - **Database**—Select **Traffic Log**.
   - **Scheduled**—Enable **Scheduled** and the report will run every night.
   - **Time Frame**—30 days
   - **Selected Columns**—Select **Source address** or **Source User** (if you have User-ID configured), which will identify the infected client host in the report, and **Destination address**, which will be the sinkhole address.
   - In the section at the bottom of the screen, create a custom query for traffic to the sinkhole address (10.15.0.20 in this example). You can either enter the destination address in the **Query Builder** window (**addr.dst in 10.15.0.20**) or select the following in each column and click **Add**: Connector = and, Attribute = Destination Address, Operator = in, and Value = 10.15.0.20. Click **Add** to add the query.

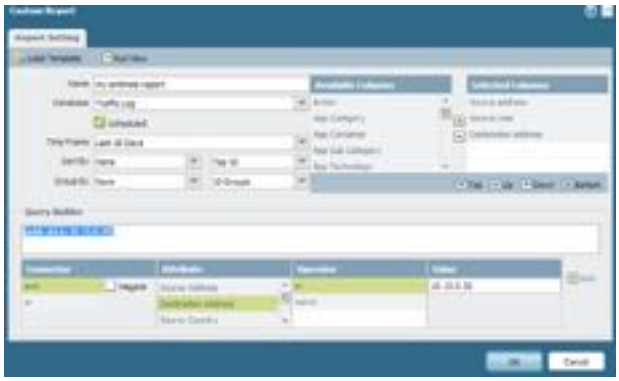

4. Click **Run Now** to run the report. The report will show all client hosts that have sent traffic to the sinkhole address, which indicates that they are most likely infected. You can now track down the hosts and check them for spyware.

5. To view scheduled reports that have run, select **Monitor > Reports**.

 © Palo Alto Networks, Inc.

# Content Delivery Network Infrastructure for Dynamic Updates

Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to the Palo Alto Networks devices. The devices access the web resources in the CDN to perform various App-ID and Content-ID functions. For enabling and scheduling the content updates, see Manage Content Updates.

The following table lists the web resources that the firewall accesses for a feature or application:

| Resource | URL | Static Addresses (If a static server is required) |
|---|---|---|
| Application Database | • updates.paloaltonetworks.com:443 | staticupdates.paloaltonetworks.com or the IP address 199.167.52.15 |
| Threat/Antivirus Database | • updates.paloaltonetworks.com:443<br>• downloads.paloaltonetworks.com:443<br>As a best practice, set the update server to updates.paloaltonetworks.com.This allows the Palo Alto Networks device to receive content updates from the server closest to it in the CDN infrastructure. | staticupdates.paloaltonetworks.com or the IP address 199.167.52.15 |
| PAN-DB URL Filtering | *.urlcloud.paloaltonetworks.com<br>Resolves to the primary URL s0000.urlcloud.paloaltonetworks.com and is then redirected to the regional server that is closest:<br>• s0100.urlcloud.paloaltonetworks.com<br>• s0200.urlcloud.paloaltonetworks.com<br>• s0300.urlcloud.paloaltonetworks.com<br>• s0500.urlcloud.paloaltonetworks.com | Static IP addresses are not available. However, you can manually resolve a URL to an IP address and allow access to the regional server IP address. |
| BrightCloud URL Filtering | • database.brightcloud.com:443/80<br>• service.brightcloud.com:80 | Contact BrightCloud Customer Support. |

© Palo Alto Networks, Inc.

| Resource | URL | Static Addresses (If a static server is required) |
|---|---|---|
| WildFire | • beta.wildfire.paloaltonetworks.com:443/80<br>• beta-s1.wildfire.paloaltonetworks.com:443/80<br> Beta sites are only accessed by a firewall running a Beta release version.<br>• mail.wildfire.paloaltonetworks.com:25<br>• wildfire.paloaltonetworks.com:443/80 | • mail.wildfire.paloaltonetworks.com:25 or the IP address 54.241.16.83<br>• wildfire.paloaltonetworks.com:443/80 or 54.241.8.199<br>The regional URL/IP addresses are as follows:<br>• ca-s1.wildfire.paloaltonetworks.com:44 or 54.241.34.71<br>• va-s1.wildfire.paloaltonetworks.com:443or 174.129.24.252<br>• eu-s1.wildfire.paloaltonetworks.com:443 or 54.246.95.247<br>• sg-s1.wildfire.paloaltonetworks.com:443or 54.251.33.241<br>• jp-s1.wildfire.paloaltonetworks.com:443 or 54.238.53.161<br>• portal3.wildfire.paloaltonetworks.com:443/80 or 54.241.8.199<br>• ca-s3.wildfire.paloaltonetworks.com:443 or 54.241.34.71<br>• va-s3.wildfire.paloaltonetworks.com:443 or 23.21.208.35<br>• eu-s3.wildfire.paloaltonetworks.com:443 or 54.246.95.247<br>• sg-s3.wildfire.paloaltonetworks.com:443 or 54.251.33.241<br>• jp-s3.wildfire.paloaltonetworks.com:443 or 54.238.53.161<br>• wildfire.paloaltonetworks.com.jp:443/80 or 180.37.183.53<br>• wf1.wildfire.paloaltonetowrks.jp:443 or 180.37.180.37<br>• wf2.wildfire.paloaltonetworks.jp:443 or 180.37.181.18<br>• portal3.wildfire.paloaltonetworks.jp:443/80 or 180.37.183.53 |


© Palo Alto Networks, Inc.

# Threat Prevention Resources

For more information on Threat Prevention, refer to the following sources:

- Creating Custom Threat Signatures
- Threat Prevention Deployment
- Understanding DoS Protection

To view a list of Threats and Applications that Palo Alto Networks products can identify, use the following links:

- Applipedia—Provides details on the applications that Palo Alto Networks can identify.
- Threat Vault—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

# Decryption

Palo Alto Networks firewalls provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Decryption on a Palo Alto Networks firewall includes the capability to enforce security policies on encrypted traffic, where otherwise the encrypted traffic might not be blocked and shaped according to your configured security settings. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted traffic. Enabling decryption on a Palo Alto Networks firewall can include preparing the keys and certificates required for decryption, creating a decryption policy, and configuring decryption port mirroring. See the following topics to learn about and configure decryption:

- Decryption Overview
- Decryption Concepts
- Configure SSL Forward Proxy
- Configure SSL Inbound Inspection
- Configure SSH Proxy
- Configure Decryption Exceptions
- Enable Users to Opt Out of SSL Decryption
- Configure Decryption Port Mirroring
- Temporarily Disable SSL Decryption

© Palo Alto Networks, Inc.

# Decryption Overview

Secure Sockets Layer (SSL) and Secure Shell (SSH) are encryption protocols used to secure traffic between two entities, such as a web server and a client. SSL and SSH encapsulate traffic, encrypting data so that it is meaningless to entities other than the client and server with the keys to decode the data and the certificates to affirm trust between the devices. Traffic that has been encrypted using the protocols SSL and SSH can be decrypted to ensure that these protocols are being used for the intended purposes only, and not to conceal unwanted activity or malicious content.

Palo Alto Networks firewalls decrypt encrypted traffic by using keys to transform strings (passwords and shared secrets) from ciphertext to plaintext (decryption) and from plaintext back to ciphertext (re-encrypting traffic as it exits the device). Certificates are used to establish the firewall as a trusted third party and to create a secure connection. SSL decryption (both forward proxy and inbound inspection) requires certificates to establish trust between two entities in order to secure an SSL/TLS connection. Certificates can also be used when excluding servers from SSL decryption. You can integrate a hardware security module (HSM) with a firewall to enable enhanced security for the private keys used in SSL forward proxy and SSL inbound inspection decryption. To learn more about storing and generating keys using an HSM and integrating an HSM with your firewall, see Secure Keys with a Hardware Security Module. SSH decryption does not require certificates.

Palo Alto Networks firewall decryption is policy-based, and can be used to decrypt, inspect, and control both inbound and outbound SSL and SSH connections. Decryption policies allow you to specify traffic for decryption according to destination, source, or URL category and in order to block or restrict the specified traffic according to your security settings. The firewall uses certificates and keys to decrypt the traffic specified by the policy to plaintext, and then enforces App-ID and security settings on the plaintext traffic, including Decryption, Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles. After traffic is decrypted and inspected on the firewall, the plaintext traffic is re-encrypted as it exits the firewall to ensure privacy and security. Use policy-based decryption on the firewall to achieve outcomes such as the following:

- Prevent malware concealed as encrypted traffic from being introduced into an corporate network.
- Prevent sensitive corporate information from moving outside the corporate network.
- Ensure the appropriate applications are running on a secure network.
- Selectively decrypt traffic; for example, exclude traffic for financial or healthcare sites from decryption by configuring a decryption exception.

The three decryption policies offered on the firewall, SSL Forward Proxy, SSL Inbound Inspection, and SSH Proxy, all provide methods to specifically target and inspect SSL outbound traffic, SSL inbound traffic, and SSH traffic, respectively. The decryption policies provide the settings for you to specify what traffic to decrypt and decryption profiles can be selected when creating a policy, in order to apply more granular security settings to decrypted traffic, such as checks for server certificates, unsupported modes, and failures. This policy-based decryption on the firewall gives you visibility into and control of SSL and SSH encrypted traffic according to configurable parameters.

You can also choose to extend a decryption configuration on the firewall to include Decryption Mirroring, which allows for decrypted traffic to be forwarded as plaintext to a third party solution for additional analysis and archiving.

 © Palo Alto Networks, Inc.

# Decryption Concepts

To learn about keys and certificates for decryption, decryption policies, and decryption port mirroring, see the following topics:

- Keys and Certificates for Decryption Policies
- SSL Forward Proxy
- SSL Inbound Inspection
- SSH Proxy
- Decryption Exceptions
- Decryption Mirroring

## Keys and Certificates for Decryption Policies

Keys are strings of numbers that are typically generated using a mathematical operation involving random numbers and large primes. Keys are used to transform other strings—such as passwords and shared secrets—from plaintext to ciphertext (called *encryption*) and from ciphertext to plaintext (called *decryption*). Keys can be symmetric (the same key is used to encrypt and decrypt) or asymmetric (one key is used for encryption and a mathematically related key is used for decryption). Any system can generate a key.

X.509 certificates are used to establish trust between a client and a server in order to establish an SSL connection. A client attempting to authenticate a server (or a server authenticating a client) knows the structure of the X.509 certificate and therefore knows how to extract identifying information about the server from fields within the certificate, such as its FQDN or IP address (called a *common name* or *CN* within the certificate) or the name of the organization, department, or user to which the certificate was issued. All certificates must be issued by a certificate authority (CA). After the CA verifies a client or server, the CA issues the certificate and signs it using its private key.

With a decryption policy configured, an SSL/TLS session between the client and the server is established only if the firewall trusts the CA that signed the server certificate. In order to establish trust, the firewall must have the server root CA certificate in its certificate trust list (CTL) and use the public key contained in that root CA certificate to verify the signature. The firewall then presents a copy of the server certificate signed by the Forward Trust certificate for the client to authenticate. You can also configure the firewall to use an enterprise CA as a forward trust certificate for SSL Forward Proxy. If the firewall does not have the server root CA certificate in its CTL, the firewall will present a copy of the server certificate signed by the Forward Untrust certificate to the client. The Forward Untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites hosted by a server with untrusted certificates.

For detailed information on certificates, see Certificate Management.

To control the trusted CAs that your device trusts, use the **Device > Certificate Management > Certificates > Default Trusted Certificate Authorities** tab on the firewall web interface.

Table: Palo Alto Networks Device Keys and Certificates describes the different keys and certificates used by Palo Alto Networks devices for decryption. As a best practice, use different keys and certificates for each usage.

**Table: Palo Alto Networks Device Keys and Certificates**

| Key/Certificate Usage | Description |
|---|---|
| Forward Trust | The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall trusts. To configure a Forward Trust certificate on the firewall, see Step 2 in the Configure SSL Forward Proxy task. By default, the firewall determines the key size to use for the client certificate based on the key size of the destination server. However, you can also set a specific key size for the firewall to use. See Configure the Key Size for SSL Forward Proxy Server Certificates. For added security, store the forward trust certificate on a Hardware Security Module (HSM), see Store Private Keys on an HSM. |
| Forward Untrust | The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall does not trust. To configure a Forward Untrust certificate on the firewall, see Step 3 in the Configure SSL Forward Proxy task. |

 © Palo Alto Networks, Inc.

| Key/Certificate Usage | Description |
|---|---|
| SSL Exclude Certificate | Certificates for servers that you want to exclude from SSL decryption. For example, if you have SSL decryption enabled, but have certain servers that you do not want included in SSL decryption, such as the web services for your HR systems, you would import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Exclude a Server from Decryption. |
| SSL Inbound Inspection | The certificate used to decrypt inbound SSL traffic for inspection and policy enforcement. For this application, you would import the server certificate for the servers for which you are performing SSL inbound inspection, or store them on an HSM (see Store Private Keys on an HSM). |

## SSL Forward Proxy

Use an SSL Forward Proxy decryption policy to decrypt and inspect SSL/TLS traffic from internal users to the web. SSL Forward Proxy decryption prevents malware concealed as SSL encrypted traffic from being introduced to your corporate network; for example, if an employee is using her Gmail account from her corporate office and opens an email attachment that contains a virus, SSL Forward Proxy decryption will prevent the virus from infecting the client system and entering the corporate network.

With SSL Forward Proxy decryption, the firewall resides between the internal client and outside server. The firewall uses Forward Trust or Forward Untrust certificates to establish itself as a trusted third party to the session between the client and the server (For details on certificates, see Keys and Certificates for Decryption Policies). When the client initiates an SSL session with the server, the firewall intercepts the client SSL request and forwards the SSL request to the server. The server sends a certificate intended for the client that is intercepted by the firewall. If the server certificate is signed by a CA that the firewall trusts, the firewall creates a copy of the server certificate signed by the Forward Trust certificate and sends the certificate to the client to authenticate. If the server certificate is signed by a CA that the firewall does not trust, the firewall creates a copy of the server certificate and signs it with the Forward Untrust certificate and sends it to the client. In this case, the client sees a block page warning that the site they're attempting to connect to is not trusted and the client can choose to proceed or terminate the session. When the client authenticates the certificate, the SSL session is established with the firewall functioning as a trusted forward proxy to the site that the client is accessing.

As the firewall continues to receive SSL traffic from the server that is destined for the client, it decrypts the SSL traffic into clear text traffic and applies security policies to the traffic. The traffic is then re-encrypted on the firewall and the firewall forwards the encrypted traffic to the client.

Figure: SSL Forward Proxy shows this process in detail.

 © Palo Alto Networks, Inc.

**Figure: SSL Forward Proxy**

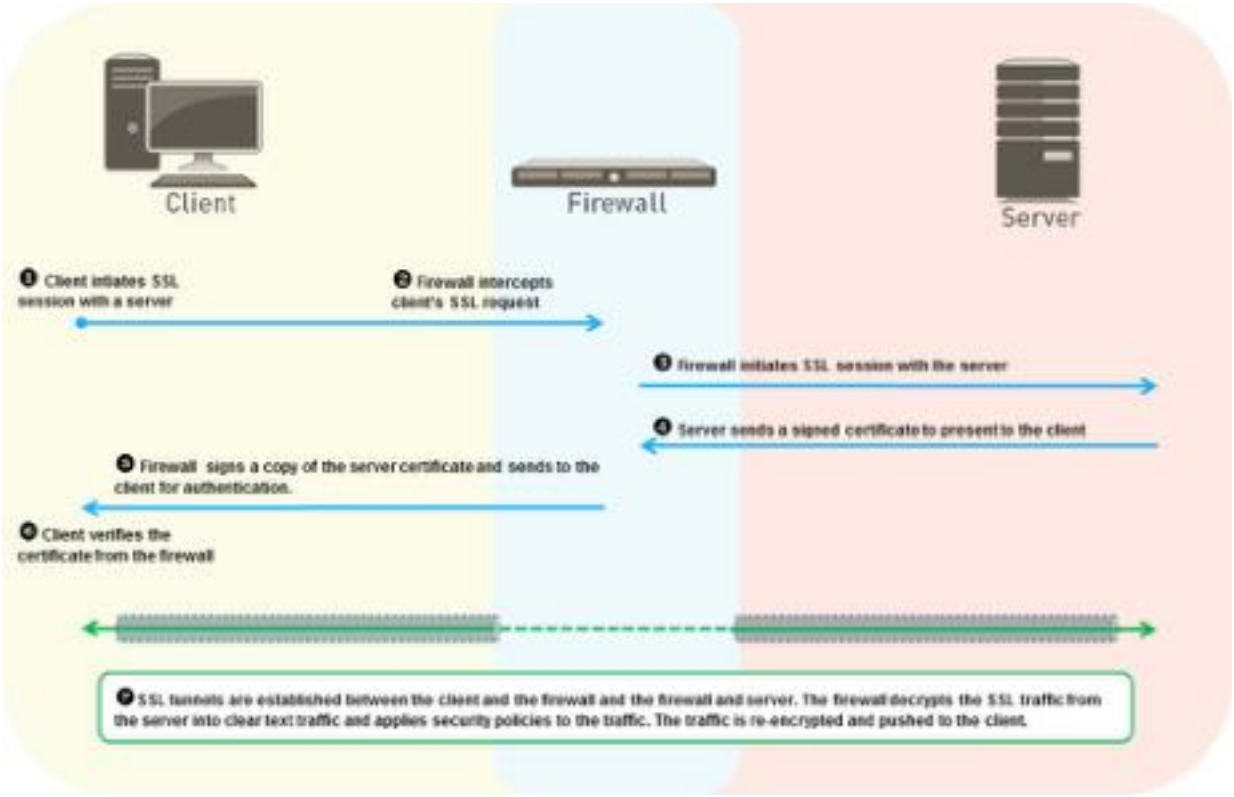


See Configure SSL Forward Proxy for details on configuring SSL Forward Proxy.

## SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL traffic from a client to a targeted server (any server you have the certificate for and can import it onto the firewall). For example, if an employee is remotely connected to a web server hosted on the company network and is attempting to add restricted internal documents to his Dropbox folder (which uses SSL for data transmission), SSL Inbound Inspection can be used to ensure that the sensitive data does not move outside the secure company network by blocking or restricting the session.

Configuring SSL Inbound Inspection includes importing the targeted server certificate and key on to the firewall. Because the targeted server certificate and key are imported on the firewall, the firewall is able to access the SSL session between the server and the client and decrypt and inspect traffic transparently, rather than functioning as a proxy. The firewall is able to apply security policies to the decrypted traffic, detecting malicious content and controlling applications running over this secure channel.

Figure: SSL Inbound Inspection shows this process in detail.

**Figure: SSL Inbound Inspection**

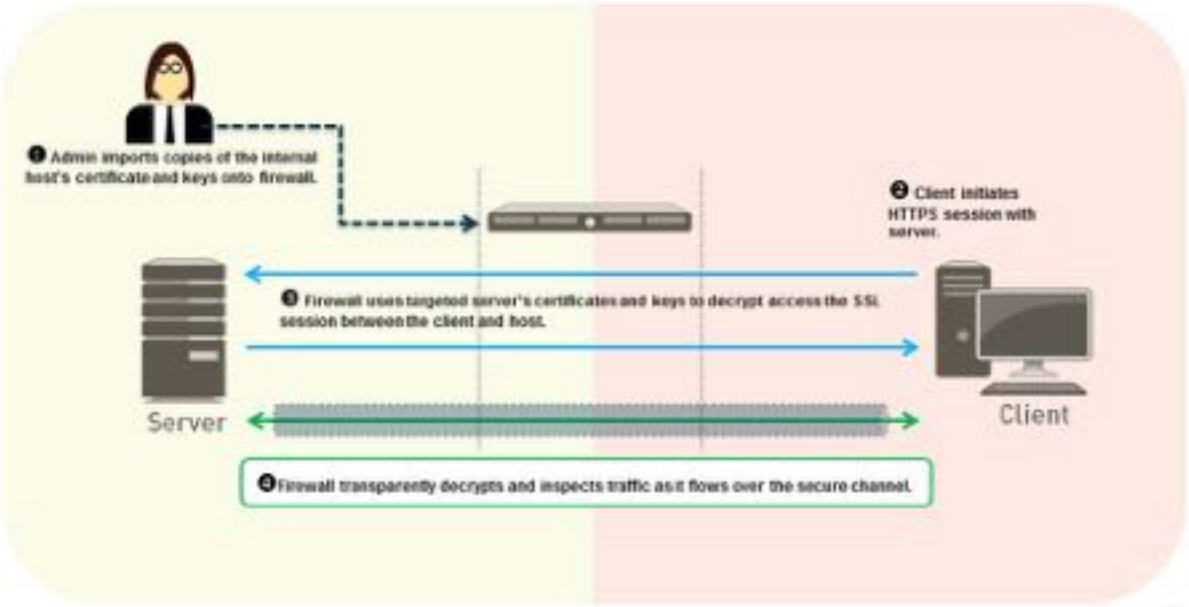


See Configure SSL Inbound Inspection for details on configuring SSL Inbound Inspection.

 © Palo Alto Networks, Inc.

## SSH Proxy

SSH Proxy provides the capability for the firewall to decrypt inbound and outbound SSH connections passing through the firewall, in order to ensure that SSH is not being used to tunnel unwanted applications and content. SSH decryption does not require any certificates and the key used for SSH decryption is automatically generated when the firewall boots up. During the boot up process, the firewall checks to see if there is an existing key. If not, a key is generated. This key is used for decrypting SSH sessions for all virtual systems configured on the device. The same key is also used for decrypting all SSH v2 sessions.

In an SSH Proxy configuration, the firewall resides between a client and a server. When the client sends an SSH request to the server, the firewall intercepts the request and forwards the SSH request to the server. The firewall then intercepts the server response and forwards the response to the client, establishing an SSH tunnel between the firewall and the client and an SSH tunnel between the firewall and the server, with firewall functioning as a proxy. As traffic flows between the client and the server, the firewall is able to distinguish whether the SSH traffic is being routed normally or if it is using SSH tunneling (port forwarding). Content and threat inspections are not performed on SSH tunnels; however, if SSH tunnels are identified by the firewall, the SSH tunneled traffic is blocked and restricted according to configured security policies.

Figure: SSH Proxy Decryption shows this process in detail.

**Figure: SSH Proxy Decryption**

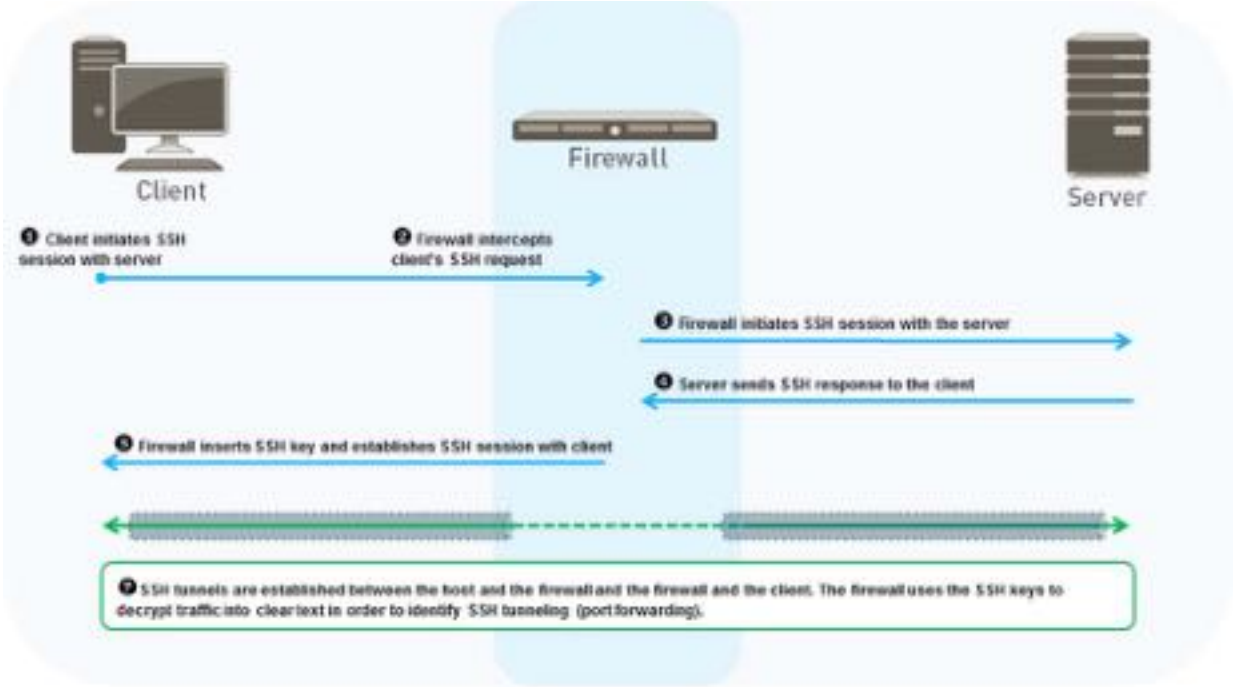


See Configure SSH Proxy for details on configuring an SSH Proxy policy.

## Decryption Exceptions

Traffic can also be excluded from decryption according to matching criteria (using a decryption policy), a targeted server traffic can be excluded from decryption (using certificates), and some applications are excluded from decryption by default.

Applications that do not function properly when decrypted by the firewall and are automatically excluded from SSL decryption. The applications that are excluded from SSL decryption by default are excluded because these applications often fail when decrypted due to the application looking for specific details in the certificate that might not be present in the certificate generated for SSL Forward Proxy. Refer to the KB article List of Applications Excluded from SSL Decryption for a current list of applications excluded by default from SSL decryption on the firewall.

You can configure decryption exceptions for certain URL categories or applications that either do not work properly with decryption enabled or for any other reason, including for legal or privacy purposes. You can use a decryption policy to exclude traffic from decryption based on source, destination, URL category, service (port or protocol), and TCP port numbers. For example, with SSL decryption enabled, you can exclude traffic that is categorized as financial or health-related from decryption using the URL category selection.

You can also exclude servers from SSL decryption based on the Common Name (CN) in the server certificate. For example, if you have SSL decryption enabled but have certain servers that you do not want included in SSL decryption, such as the web services for your HR systems, you can exclude those servers from decryption by importing the server certificate onto the firewall and modifying the certificate to be an **SSL Exclude Certificate**.

To exclude traffic from decryption based on application, source, destination, URL category or to exclude a specific server from decryption, see Configure Decryption Exceptions.

 © Palo Alto Networks, Inc.

## Decryption Mirroring

The decryption mirroring feature provides the capability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures–such as NetWitness or Solera–for archiving and analysis. This feature is necessary for organizations that require comprehensive data capture for forensic and historical purposes or data leak prevention (DLP) functionality. Decryption mirroring is available on PA-7000 Series, PA-5000 Series and PA-3000 Series platforms only and requires that a free license be installed to enable this feature.

Keep in mind that the decryption, storage, inspection, and/or use of SSL traffic is governed in certain countries and user consent might be required in order to use the decryption mirror feature. Additionally, use of this feature could enable malicious users with administrative access to the firewall to harvest usernames, passwords, social security numbers, credit card numbers, or other sensitive information submitted using an encrypted channel. Palo Alto Networks recommends that you consult with your corporate council before activating and using this feature in a production environment.

Figure: Decryption Port Mirroring shows the process for mirroring decrypted traffic and the section Configure Decryption Port Mirroring describes how to license and enable this feature.

**Figure: Decryption Port Mirroring**

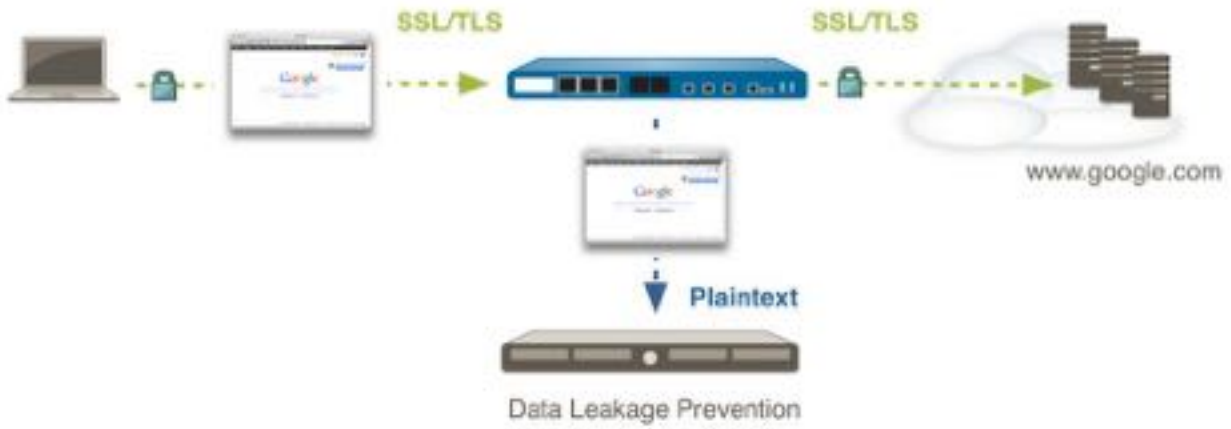

# Configure SSL Forward Proxy

Configuring SSL Forward Proxy decryption on the firewall requires setting up the certificates needed for SSL Forward Proxy decryption and creating an SSL Forward Proxy decryption policy. The firewall can use self-signed certificates or certificates signed by an enterprise CA to perform SSL Forward Proxy decryption.

By default, the firewall determines the key size to use for the client certificates it generates based on the key size of the destination server certificate. You can optionally set a static key size to use regardless of the key size of the destination server certificate. See Configure the Key Size for SSL Forward Proxy Server Certificates.

Use the following task to configure SSL Forward Proxy, including how to set up the certificates and create a decryption policy.

| Configure SSL Forward Proxy | | |
|---|---|---|
| Step 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. | View configured interfaces on the **Network > Interfaces > Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is. |


© Palo Alto Networks, Inc.

**Configure SSL Forward Proxy**

Step 2 Configure the forward trust certificate. Use either a self-signed certificate or a certificate signed by an enterprise CA.

**Using self-signed certificates**

When the certificate of the server that the client is connecting to is signed by a CA that is on the firewall trusted CA list, the firewall signs a copy of the server certificate with a self-signed forward trust certificate to present to the client for authentication. In this case, the self-signed certificate must be imported onto each client system so that the client recognizes the firewall as a trusted CA.

Use self-signed certificates for SSL Forward Proxy decryption if you do not use an enterprise CA or if you are only intended to perform decryption for a limited number of client systems (or if you are planning to use a centralized deployment).

**To use a self-signed certificate:**

1. Select **Device > Certificate Management > Certificates**.
2. Click **Generate** at the bottom of the window.
3. Enter a **Certificate Name**, such as *my-fwd-trust*.
4. Enter a **Common Name**, such as **192.168.2.1**. This should be the IP or FQDN that will appear in the certificate. In this case, we are using the IP of the trust interface. Avoid using spaces in this field.
5. Leave the **Signed By** field blank.
6. Click the **Certificate Authority** check box to enable the firewall to issue the certificate. Selecting this check box creates a certificate authority (CA) on the firewall that is imported to the client browsers, so clients trust the firewall as a CA.
7. Click **Generate** to generate the certificate.
8. Click the new certificate *my-fwd-trust* to modify it and enable the **Forward Trust Certificate** option.
9. Export the forward trust certificate for import into client systems by highlighting the certificate and clicking **Export** at the bottom of the window. Choose PEM format, and do not select the **Export private key** option. Because the certificate is self-signed, import it into the browser trusted root CA list on the client systems in order for the clients to trust it. When importing to the client browser, ensure the certificate is added to the Trusted Root Certification Authorities certificate store. On Windows systems, the default import location is the Personal certificate store. You can also simplify this process by using a centralized deployment, such as an Active Directory Group Policy Object (GPO).

   If the forward trust certificate is not imported on the client systems, users will see certificate warnings for each SSL site they visit.
10. Click **OK** to save.

## Configure SSL Forward Proxy

### Using an Enterprise CA

An enterprise CA can issue a signing certificate which the firewall can use to then sign the certificates for sites requiring SSL decryption. Send a Certificate Signing Request (CSR) for the enterprise CA to sign and validate. The firewall can then use the signed enterprise CA certificate for SSL Forward Proxy decryption. Because the enterprise CA is already trusted by the client systems, with this option, you do not need to distribute the certificate to client systems prior to configuring decryption.

**To use an enterprise CA signed certificate, generate a CSR:**

1. Select **Device > Certificate Management > Certificates** and click **Generate**.
2. Enter a **Certificate Name**, such as *my-fwd-proxy*.
3. In the **Signed By** drop-down, select **External Authority (CSR)**.
4. (Optional) If your enterprise CA requires it, add **Certificate Attributes** to further identify the firewall details, such as Country or Department.
5. Click **OK** to save the CSR. The pending certificate is now displayed on the **Device Certificates** tab.
6. Export the CSR:
   a. Select the pending certificate displayed on the **Device Certificates** tab.
   b. Click **Export** to download and save the certificate file.

   Leave **Export private key** unselected in order to ensure that the private key remains securely on the firewall.
   c. Click **OK**.
   d. Provide the certificate file to your enterprise CA. When you receive the signed enterprise CA certificate from your enterprise CA, save the signed enterprise CA certificate for import onto the firewall.
7. Import the signed enterprise CA onto the firewall:
   a. Select **Device > Certificate Management > Certificates** and click **Import**.
   b. Enter the pending **Certificate Name** exactly (in this case, *my-fwd-trust*). The **Certificate Name** that you enter must exactly match the pending certificate name in order for the pending certificate to be validated.
   c. Select the signed **Certificate File** that you received from your enterprise CA.
   d. Click **OK**. The certificate is displayed as valid with the Key and CA check boxes selected.
   e. Select the validated certificate, in this case, *my-fwd-proxy*, to enable it as a **Forward Trust Certificate** to be used for SSL Forward Proxy decryption.
   f. Click **OK**.


© Palo Alto Networks, Inc.

| Configure SSL Forward Proxy | | |
|---|---|---|
| Step 3 | Configure the forward untrust certificate.<br><br>With SSL Forward Proxy decryption, when the site the client is connecting to uses a certificate signed by a CA that is not in the firewall trusted CA list, the firewall presents a forward untrust certificate to the client. The forward untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites with untrusted certificates. | 1. Click **Generate** at the bottom of the certificates page.<br>2. Enter a **Certificate Name**, such as *my-fwd-untrust*.<br>3. Set the **Common Name**, for example **192.168.2.1**. Leave **Signed By** blank.<br>4. Click the **Certificate Authority** check box to enable the firewall to issue the certificate.<br>5. Click **Generate** to generate the certificate.<br>6. Click **OK** to save.<br>7. Click the new *my-ssl-fw-untrust* certificate to modify it and enable the **Forward Untrust Certificate** option.<br><br> Do not export the forward untrust certificate for import into client systems. If the forward trust certificate is imported on client systems, the users will not see certificate warnings for SSL sites with untrusted certificates.<br><br>8. Click **OK** to save. |
| Step 4 | (Optional) Set the key size of the SSL Forward Proxy certificates that the firewall presents to clients.<br><br> Changing the key size clears the current certificate cache. | 1. Select **Device > Setup > Session** and, in the Decryption Settings section, click **Forward Proxy Server Certificate Settings**.<br>2. Select a **Key Size**: **Defined by destination host** (default), **1024-bit RSA**, or **2048-bit RSA**. |
| Step 5 | (Optional) Create a decryption profile.<br><br>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. An SSL Forward Proxy decryption profile can be used to:<br>• Perform checks for server certificates, unsupported modes, and failures.<br>• Block or restrict traffic.<br>• Enforce protocol versions, strong cipher suites, and certificate validation. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Select **SSL Decryption > SSL Forward Proxy** to block and control specific aspects of SSL tunneled traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting **Block sessions if resources not available**.<br>3. Click **OK** to save the profile. |

| Configure SSL Forward Proxy | |
|---|---|
| Step 6 Configure a decryption policy. | 1. Select **Policies > Decryption** and click **Add**.<br>2. On the **General** tab, give the policy a descriptive **Name**.<br>3. On the **Source** and **Destination** tabs, select **Any** for the **Source Zone** and **Destination Zone** to decrypt all SSL traffic destined for an external server. If you want to specify traffic from or to certain sources or destinations for decryption, click **Add**.<br>4. On the **URL Category** tab, select **Any** to decrypt all traffic. Alternatively, **Add** URL categories to apply the profile to only those websites that fall under the selected categories.<br> The option to decrypt traffic based on URL category is also useful when excluding certain sites (such as financial or healthcare-related sites ) from decryption. See Configure Decryption Exceptions.<br>5. On the **Options** tab, select **Decrypt** and select **SSL Forward Proxy** as the **Type** of decryption to perform.<br>6. (Optional) Select a **Decryption Profile** to apply additional settings to decrypted traffic (see Step 5).<br>7. Click **OK** to save. |
| Step 7 **Commit** the configuration. | With an SSL Forward Proxy decryption policy enabled, all traffic identified by the policy is decrypted. Decrypted traffic is blocked and restricted according to the profiles configured on the firewall (including the decryption profiles associated with the policy and Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles). Traffic is re-encrypted as it exits the firewall. |


© Palo Alto Networks, Inc.

# Configure SSL Inbound Inspection

Configuring SSL Inbound Inspection includes installing the targeted server certificate on the firewall and creating an SSL Inbound Inspection decryption policy.

Use the following task to configure SSL Inbound Inspection.

| Configure SSL Inbound Inspection | | |
|---|---|---|
| Step 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. | View configured interfaces on the **Network > Interfaces > Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is. |
| Step 2 | Ensure that the targeted server certificate is installed on the firewall. | On the web interface, select **Device > Certificate Management > Certificates > Device Certificates** to view certificates installed on the firewall.<br>To import the targeted server certificate onto the firewall:<br>1. On the **Device Certificates** tab, select **Import**.<br>2. Enter a descriptive **Certificate Name**.<br>3. Browse for and select the targeted server **Certificate File**.<br>4. Click **OK**. |
| Step 3 | (Optional) Create a Decryption profile.<br>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. An SSL Inbound Inspection decryption profile can be used to perform checks for unsupported protocol versions, unsupported cipher suites, and failures. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Select **SSL Decryption > Inbound Inspection** to block and control specific aspects of inbound SSL traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting **Block sessions if resources not available**.<br>3. Click **OK t**o save the profile. |

| Configure SSL Inbound Inspection | |
|---|---|
| Step 4 Configure a decryption policy. | 1. Select **Policies > Decryption** and click **Add**.<br>2. On the **General** tab, give the policy a descriptive **Name**.<br>3. On the **Destination** tab, **Add** the **Destination Address** of the targeted server.<br>4. On the **URL Category** tab, select **Any** to decrypt all traffic. Alternatively, **Add** URL categories to apply the profile to only those websites that fall under the selected categories.<br> The option to decrypt traffic based on URL category is also useful when excluding certain sites (such as financial or healthcare-related sites ) from decryption. See Configure Decryption Exceptions.<br>5. On the **Options** tab, select **Decrypt** and select **SSL Inbound Inspection** as the **Type** of traffic to decrypt.<br>Select the **Certificate** for the internal server that is the destination of the inbound SSL traffic.<br>6. (Optional) Select a **Decryption Profile** to apply additional settings to decrypted traffic.<br>7. Click **OK** to save. |
| Step 5 **Commit** the configuration. | With an SSL Inbound Inspection decryption policy enabled, all SSL traffic identified by the policy is decrypted and inspected. Decrypted traffic is blocked and restricted according to the profiles configured on the firewall (including the decryption profiles associated with the policy and Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles). Traffic is re-encrypted as it exits the firewall. |


© Palo Alto Networks, Inc.

# Configure SSH Proxy

Configuring SSH Proxy does not require certificates and the key used to decrypt SSH sessions is generated automatically on the firewall during boot up.

Use the following task to configure SSH Proxy decryption.

| Configure SSH Proxy Decryption | | |
|---|---|---|
| Step 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. Decryption can only be performed on virtual wire, Layer 2, or Layer 3 interfaces. | View configured interfaces on the **Network > Interfaces > Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is. |
| Step 2 | (Optional) Create a Decryption profile.<br>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. An SSH proxy decryption profile can be used to perform checks for unsupported protocol versions, unsupported cipher suites, and for failures. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Select the **SSH** tab to block and control specific aspects of SSH tunneled traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting **Block sessions if resources not available**.<br>3. Click **OK** to save the profile. |
| Step 3 | Configure a decryption policy. | 1. Select **Policies > Decryption** and click **Add**.<br>2. On the **General** tab, give the policy a descriptive **Name**.<br>3. On the **Source** and **Destination** tabs, select **Any** to decrypt all SSH traffic.<br>4. On the **URL Category** tab, select **Any** to decrypt all SSH traffic.<br>5. On the **Options** tab, select **Decrypt** and select **SSH Proxy** as the **Type** of traffic to decrypt.<br>6. (Optional) Select a **Decryption Profile** to apply additional settings to decrypted traffic.<br>7. Click **OK** to save. |
| Step 4 | **Commit** the configuration. | With the an SSH Proxy decryption policy enabled, all SSH traffic identified by the policy is decrypted and identified as either regular SSH traffic or as SSH tunneled traffic. SSH tunneled traffic is blocked and restricted according to the profiles configured on the firewall. Traffic is re-encrypted as it exits the firewall. |

# Configure Decryption Exceptions

You can purposefully exclude traffic from decryption based on source, destination, URL category, and service (ports and protocols). You can also exclude a specific server from decryption. See the following topics to configure Decryption Exceptions:

- Exclude Traffic from Decryption
- Exclude a Server from Decryption

## Exclude Traffic from Decryption

To exclude applications or certain traffic from decryption, create a decryption policy rule that defines the traffic to be excluded from decryption and set the policy action to **No Decrypt**. Exclude traffic from decryption based on source, destination, URL category, and service (ports and protocols). Because policy rules are compared against incoming traffic in sequence, make sure that a decryption exclusion rule is listed first in your decryption policy.

| **Exclude Traffic from a Decryption Policy** | | |
|---|---|---|
| Step 1 | (Optional) Enable the firewall to block sessions with expired certificates and/or untrusted certificate issuers, without decrypting those sessions. | 1. Select **Objects > Decryption Profile** and click **Add**.<br>2. Give the profile rule a descriptive **Name** such as: Block certs and issues for No Decrypt traffic.<br>3. Select **No Decryption** to block and control specific aspects of traffic that you are excluding from decryption. Select **Block sessions with expired certificates** and/or **Block sessions with untrusted issuers**.<br>4. Click **OK t**o save the profile rule. |
| Step 2 | Configure a decryption policy rule.<br><br>Use a decryption policy rule to exclude traffic from decryption based on source and destination zones and addresses, URL categories, ports, and protocols. This example shows how to exclude traffic categorized as financial or health-related from SSL Forward Proxy decryption. | 1. Go to **Policies > Decryption** and click **Add**.<br>2. Give the rule a descriptive **Name**, such as *No-Decrypt-Finance-Health*.<br>3. On the **Source** and **Destination** tabs, select **Any** for the **Source Zone** and **Destination Zone** to apply the *No-Decrypt-Finance-Health* rule to all SSL traffic destined for an external server.<br>4. On the **URL Category** tab, **Add** the URL categories financial-services and health-and-medicine to the policy, specifying that traffic that matches these categories will not be decrypted.<br>5. On the **Options** tab, select **No Decrypt**.<br>6. (Optional) Attach the decryption profile created in Step 1 to the policy. A decryption profile can be attached to a No Decrypt policy, to block and control aspects of the traffic being excluded from decryption.<br>7. Click **OK** to save the *No-Decrypt-Finance-Health* rule. |

 © Palo Alto Networks, Inc.

| Exclude Traffic from a Decryption Policy | | |
|---|---|---|
| Step 3 | Move the decryption exclusion rule to the to the top of your decryption policy. | On the **Decryption > Policies** page, select the policy *No-Decrypt-Finance-Health*, and click **Move Up** until it appears at the top of the list (or you can drag and drop the rule).<br>Decryption rules are enforced against incoming traffic in sequence and the first rule to match to traffic is enforced—moving the **No Decrypt** rule to the top of the rule list ensures that traffic defined to be excluded from decryption remains encrypted, even if the traffic is also matched to criteria defined in other decryption rules. |
| Step 4 | Save the changes. | **Commit** the configuration. |

## Exclude a Server from Decryption

You can exclude targeted server traffic from SSL decryption based on the common name (CN) in the server certificate. For example, if you have SSL decryption enabled, you could configure a decryption exception for the server on your corporate network that hosts the web services for your HR systems.

| Exclude a Server from Decryption | |
|---|---|
| Step 1 | Import the targeted server certificate onto the firewall:<br>1. On the **Device > Certificate Management > Certificates > Device Certificates** tab, select **Import**.<br>2. Enter a descriptive **Certificate Name**.<br>3. Browse for and select the targeted server **Certificate File**.<br>4. Click **OK**. |
| Step 2 | Select the targeted server certificate on the **Device Certificates** tab and enable it as an SSL Exclude Certificate.<br>With the targeted server certificate imported on the firewall and designated as an SSL Exclude Certificate, the server traffic is not decrypted as it passes through the firewall. |

# Enable Users to Opt Out of SSL Decryption

In some cases, you might need to alert your users to the fact that the firewall is decrypting certain web traffic and allow them to terminate sessions that they do not want inspected. With SSL Opt Out enabled, the first time a user attempts to browse to an HTTPS site or application that matches your decryption policy, the firewall displays a response page notifying the user that it will decrypt the session. Users can either click **Yes** to allow decryption and continue to the site or click **No** to opt out of decryption and terminate the session. The choice to allow decryption applies to all HTTPS sites that users try to access for the next 24 hours, after which the firewall redisplays the response page. Users who opt out of SSL decryption cannot access the requested web page, or any other HTTPS site, for the next minute. After the minute elapses, the firewall redisplays the response page the next time the users attempt to access an HTTPS site.

The firewall includes a predefined SSL Decryption Opt-out Page that you can enable. You can optionally customize the page with your own text and/or images.

| **Enable Users to Opt Out of SSL Decryption** | |
|---|---|
| Step 1 (Optional) Customize the SSL Decryption Opt-out Page. | 1. Select **Device > Response Pages**.<br>2. Select the **SSL Decryption Opt-out Page** link.<br>3. Select the **Predefined** page and click **Export**.<br>4. Using the HTML text editor of your choice, edit the page.<br>5. If you want to add an image, host the image on a web server that is accessible from your end user systems.<br>6. Add a line to the HTML to point to the image. For example:<br>`<img src="http://cdn.slidesharecdn.com/Acme-logo-96x96.jpg?1382722588"/>`<br>7. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding.<br>8. Back on the firewall, select **Device > Response Pages**.<br>9. Select the **SSL Decryption Opt-out Page** link.<br>10. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.<br>11. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select shared to make it available to all virtual systems.<br>12. Click **OK** to import the file.<br>13. Select the response page you just imported and click **Close**. |
| Step 2 Enable SSL Decryption Opt Out. | 1. On the **Device > Response Pages** page, click the **Disabled** link.<br>2. Select the **Enable SSL Opt-out Page** and click **OK**.<br> <br>3. **Commit** the changes. |

 © Palo Alto Networks, Inc.

| Enable Users to Opt Out of SSL Decryption | | |
|---|---|---|
| Step 3 | Verify that the Opt Out page displays when you attempt to browse to a site. | From a browser, go to an encrypted site that matches your decryption policy.<br>Verify that the SSL Decryption Opt-out response page displays.<br> |

# Configure Decryption Port Mirroring

Before you can enable Decryption Mirroring, you must obtain and install a Decryption Port Mirror license. The license is free of charge and can be activated through the support portal as described in the following procedure. After you install the Decryption Port Mirror license and reboot the firewall, you can enable decryption port mirroring.

| Configure Decryption Port Mirroring | |
|---|---|
| Step 1 Request a license for each device on which you want to enable decryption port mirroring. | 1. Log in to the Palo Alto Networks Support site and navigate to the **Assets** tab.<br>2. Select the device entry for the device you want to license and select **Actions**.<br>3. Select **Decryption Port Mirror**. A legal notice displays.<br>4. If you are clear about the potential legal implications and requirements, click **I understand and wish to proceed**.<br>5. Click **Activate**.<br> 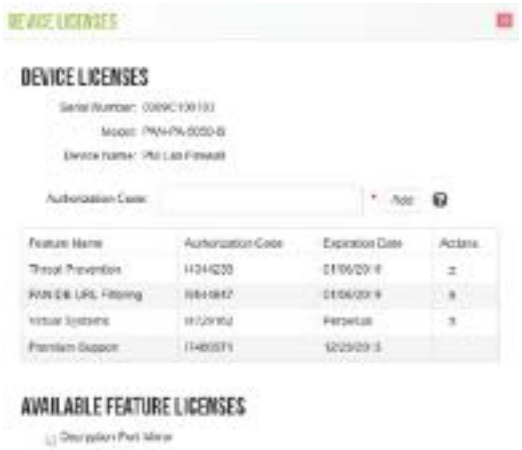 |
| Step 2 Install the Decryption Port Mirror license on the firewall. | 1. From the firewall web interface, select **Device > Licenses**.<br>2. Click **Retrieve license keys from license server**.<br>3. Verify that the license has been activated on the firewall.<br> <br>4. Reboot the firewall (**Device > Setup > Operations**). This feature will not be available for configuration until PAN-OS reloads. |

 © Palo Alto Networks, Inc.

| Configure Decryption Port Mirroring (Continued) | | |
|---|---|---|
| Step 3 | Enable the firewall to forward decrypted traffic. Superuser permission is required to perform this step. | **On a firewall with a single virtual system:**<br>1. Select **Device > Setup > Content - ID**.<br>2. Select the **Allow forwarding of decrypted content** check box.<br>3. Click **OK** to save.<br>**On a firewall with multiple virtual systems:**<br>1. Select **Device > Virtual System.**<br>2. Select a Virtual System to edit or create a new Virtual System by selecting **Add**.<br>3. Select the **Allow forwarding of decrypted content** check box.<br>4. Click **OK** to save. |
| Step 4 | Enable an Ethernet interface to be used for decryption mirroring. | 1. Select **Network > Interfaces > Ethernet**.<br>2. Select the Ethernet interface that you want to configure for decryption port mirroring.<br>3. Select **Decrypt Mirror** as the **Interface Type**.<br>This interface type will appear only if the Decryption Port Mirror license is installed.<br>4. Click **OK** to save. |
| Step 5 | Enable mirroring of decrypted traffic. | 1. Select **Objects > Decryption Profile**.<br>2. Select an **Interface** to be used for **Decryption Mirroring**.<br>The **Interface** drop-down contains all Ethernet interfaces that have been defined as the type: **Decrypt Mirror**.<br>3. Specify whether to mirror decrypted traffic before or after policy enforcement.<br>By default, the firewall will mirror all decrypted traffic to the interface before security policies lookup, which allows you to replay events and analyze traffic that generates a threat or triggers a drop action. If you want to only mirror decrypted traffic after security policy enforcement, select the **Forwarded Only** check box. With this option, only traffic that is forwarded through the firewall is mirrored. This option is useful if you are forwarding the decrypted traffic to other threat detection devices, such as a DLP device or another intrusion prevention system (IPS).<br>4. Click **OK** to save the decryption profile. |
| Step 6 | Attach the decryption profile rule (with decryption port mirroring enabled) to a decryption policy rule. All traffic decrypted based on the policy rule is mirrored. | 1. Select **Policies > Decryption**.<br>2. Click **Add** to configure a decryption policy or select an existing decryption policy to edit.<br>3. In the **Options** tab, select **Decrypt** and the **Decryption Profile** created in Step 4.<br>4. Click **OK** to save the policy. |
| Step 7 | Save the configuration. | Click **Commit**. |

# Temporarily Disable SSL Decryption

In some cases you may want to temporarily disable SSL decryption. For example, if your users are having problems accessing an encrypted site or application, you may want to disable SSL decryption in order to troubleshoot the issue. Although you could disable the associated decryption policies, modifying the policies is a configuration change that requires a Commit. Instead, use the following command to temporarily disable SSL decryption and then re-enable it after you finish troubleshooting. This command does not require a commit and it does not persist in your configuration after a reboot.

| **Temporarily Disable SSL Decryption** | |
|---|---|
| • Disable SSL Decryption | `set system setting ssl-decrypt skip-ssl-decrypt yes` |
| • Re-enable SSL Decryption | `set system setting ssl-decrypt skip-ssl-decrypt no` |

 © Palo Alto Networks, Inc.

# URL Filtering

The Palo Alto Networks URL filtering solution allows you to monitor and control how users access the web over HTTP and HTTPS.


- URL Filtering Overview
- URL Filtering Concepts
- PAN-DB Categorization
- Enable a URL Filtering Vendor
- Determine URL Filtering Policy Requirements
- Monitor Web Activity
- Configure URL Filtering
- Customize the URL Filtering Response Pages
- Configure URL Admin Override
- Enable Safe Search Enforcement
- Set Up the PAN-DB Private Cloud
- URL Filtering Use Case Examples
- Troubleshoot URL Filtering


© Palo Alto Networks, Inc.

# URL Filtering Overview

The Palo Alto Networks URL filtering solution compliments App-ID by enabling you to configure the firewall to identify and control access to web (HTTP and HTTPS) traffic and to protect your network from attack.

With URL Filtering enabled, all web traffic is compared against the URL filtering database, which contains a listing of millions of websites that have been categorized into approximately 60-80 categories. You can use these URL categories as a match criteria in policies (Captive Portal, Decryption, Security, and QoS) or attach them as URL filtering profiles in security policy, to safely enable web access and control the traffic that traverses your network.

Although the Palo Alto Networks URL filtering solution supports both BrightCloud and PAN-DB, only the PAN-DB URL filtering solution allows you to choose between the *PAN-DB Public Cloud* and the *PAN-DB Private Cloud*. Use the public cloud solution if the Palo Alto Networks next-generation firewalls on your network can directly access the Internet. If the network security requirements in your enterprise prohibit the firewalls from directly accessing the Internet, you can deploy a PAN-DB private cloud on one or more M-500 appliances that function as PAN-DB servers within your network.

- URL Filtering Vendors
- Interaction Between App-ID and URL Categories
- PAN-DB Private Cloud

 © Palo Alto Networks, Inc.

## URL Filtering Vendors

Palo Alto Networks firewalls support two URL filtering vendors:

- **PAN-DB**—A Palo Alto Networks developed URL filtering database that is tightly integrated into PAN-OS and the Palo Alto Networks threat intelligence cloud. PAN-DB provides high-performance local caching for maximum inline performance on URL lookups, and offers coverage against malicious URLs and IP addresses. As WildFire, which is a part of the Palo Alto Networks threat intelligence cloud, identifies unknown malware, zero-day exploits, and advanced persistent threats (APTs), the PAN-DB database is updated with information on malicious URLs so that you can block malware downloads, and disable Command and Control (C&C) communications to protect your network from cyber threats.

  To view a list of PAN-DB URL filtering categories, refer to https://urlfiltering.paloaltonetworks.com/CategoryList.aspx.

- **BrightCloud**—A third-party URL database that is owned by Webroot, Inc. and is integrated into PAN-OS firewalls. For information on the BrightCloud URL database, visit http://brightcloud.com.

For instructions on configuring the firewall to use one of the supported URL Filtering vendors, see Enable a URL Filtering Vendor.

## Interaction Between App-ID and URL Categories

The Palo Alto Networks URL filtering solution in combination with App-ID provides unprecedented protection against a full spectrum of cyber attacks, legal, regulatory, productivity, and resource utilization risks. While App-ID gives you control over what applications users can access, URL filtering provides control over related web activity. When combined with User-ID, you can enforce controls based on users and groups.

With today's application landscape and the way many applications use HTTP and HTTPS, you will need to use App-ID, URL filtering, or both in order to define comprehensive web access policies. App-ID signatures are granular and they allow you to identify shifts from one web-based application to another; URL filtering allows you to enforce actions based on a specific website or URL category. For example, while you can use URL filtering to control access to Facebook and/or LinkedIn, URL filtering cannot block the use of related applications such as email, chat, or other any new applications that are introduced after you implement policy. When combined with App-ID, you can control the use of related applications because of the granular application signatures that can identify each application and regulate access to Facebook while blocking access to Facebook chat, when defined in policy.

You can also use URL categories as a match criteria in policies. Instead of creating policies limited to either allow all or block all behavior, URL as a match criteria permits exception-based behavior and gives you more granular policy enforcement capabilities. For example, deny access to malware and hacking sites for all users, but allow access to users that belong to the IT-security group.

For some examples, see URL Filtering Use Case Examples.

 © Palo Alto Networks, Inc.

## PAN-DB Private Cloud

The PAN-DB private cloud is an on-premise solution that is suitable for organizations that prohibit or restrict the use of the PAN-DB public cloud service. With this on-premise solution, you can deploy one or more M-500 appliances as PAN-DB servers within your network or data center. The firewalls query the PAN-DB private cloud to perform URL lookups, instead of accessing the PAN-DB public cloud.

The process for performing URL lookups, in both the private and the public cloud is the same for the firewalls on the network. By default, the firewall is configured to access the public PAN-DB cloud. If you deploy a PAN-DB private cloud, you must configure the firewalls with a list of IP addresses or FQDNs to access the server(s) in the private cloud.

Firewalls running PAN-OS 5.0 or later versions can communicate with the PAN-DB private cloud.

When you Set Up the PAN-DB Private Cloud, you can either configure the M-500 appliance(s) to have direct Internet access or keep it completely offline. Because the M-500 appliance requires database and content updates to perform URL lookups, if the appliance does not have an active Internet connection, you must manually download the updates to a server on your network and then, import the updates using SCP into each M-500 appliance in the PAN-DB private cloud. In addition, the appliances must be able to obtain the seed database and any other regular or critical content updates for the firewalls that it services.

To authenticate the firewalls that connect to the PAN-DB private cloud, a set of default server certificates are packaged with the appliance; you cannot import or use another server certificate for authenticating the firewalls. If you change the hostname on the M-500 appliance, the appliance automatically generates a new set of certificates to authenticate the firewalls.

- M-500 Appliance for PAN-DB Private Cloud
- Differences Between the PAN-DB Public Cloud and PAN-DB Private Cloud

### M-500 Appliance for PAN-DB Private Cloud

To deploy a PAN-DB private cloud, you need one or more M-500 appliances. The M-500 appliance ships in Panorama mode, and to be deployed as PAN-DB private cloud you must set it up to operate in PAN-URL-DB mode. In the PAN-URL-DB mode, the appliance provides URL categorization services for enterprises that do not want to use the PAN-DB public cloud.

The M-500 appliance when deployed as a PAN-DB private cloud uses two ports- MGT (Eth0) and Eth1; Eth2 is not available for use. The management port is used for administrative access to the appliance and for obtaining the latest content updates from the PAN-DB public cloud or from a server on your network. For communication between the PAN-DB private cloud and the firewalls on the network, you can use the MGT port or Eth1.

The M-100 appliance cannot be deployed as a PAN-DB private cloud.

The M-500 appliance in PAN-URL-DB mode:

- Does not have a web interface, it only supports a command-line interface (CLI).

- Cannot be managed by Panorama.
- Cannot be deployed in a high availability pair.
- Does not require a URL Filtering license. The firewalls, must have a valid PAN-DB URL Filtering license to connect with and query the PAN-DB private cloud.
- Ships with a set of default server certificates that are used to authenticate the firewalls that connect to the PAN-DB private cloud. You cannot import or use another server certificate for authenticating the firewalls. If you change the hostname on the M-500 appliance, the appliance automatically generates a new set of certificates to authenticate the firewalls that it services.
- Can be reset to Panorama mode only. If you want to deploy the appliance as a dedicated Log Collector, switch to Panorama mode and then set it in log collector mode.

## Differences Between the PAN-DB Public Cloud and PAN-DB Private Cloud

| Differences | PAN-DB Public Cloud | PAN-DB Private Cloud |
|---|---|---|
| Content and Database Updates | Content (regular and critical) updates and full database updates are published multiple times during the day. The firewall checks for critical updates whenever it queries the cloud servers for URL lookups. | Content updates and full URL database updates are available once a day during the work week. |
| URL Categorization Requests | Submit URL categorization change requests using the following options:<br>• Palo Alto Networks Test A Site website.<br>• URL filtering profile setup page on the firewall.<br>• URL filtering log on the firewall. | Submit URL categorization change requests only using the Palo Alto Networks Test A Site website. |
| Unresolved URL Queries | If the firewall cannot resolve a URL query, the request is sent to the servers in the public cloud. | If the firewall cannot resolve a query, the request is sent to the M-500 appliance(s) in the PAN-DB private cloud. If there is no match for the URL, the PAN-DB private cloud sends a category *unknown* response to the firewall; the request is not sent to the public cloud unless you have configured the M-500 appliance to access the PAN-DB public cloud.<br>If the M-500 appliance(s) that constitute your PAN-DB private cloud is configured to be completely offline, it does not send any data or analytics to the public cloud. |

 © Palo Alto Networks, Inc.

# URL Filtering Concepts


- URL Categories
- URL Filtering Profile
- URL Filtering Profile Actions
- Block and Allow Lists
- Safe Search Enforcement
- Container Pages
- HTTP Header Logging
- URL Filtering Response Pages
- URL Category as Policy Match Criteria


© Palo Alto Networks, Inc.

## URL Categories

Each website defined in the URL filtering database is assigned one of approximately 60 different URL categories. There are two ways to make use of URL categorization on the firewall:

- **Block or allow traffic based on URL category**—You can create a URL Filtering profile that specifies an action for each URL category and attach the profile to a policy. Traffic that matches the policy would then be subject to the URL filtering settings in the profile. For example, to block all gaming websites you would set the block action for the URL category *games* in the URL profile and attach it to the security policy rule(s) that allow web access. See Configure URL Filtering for more information.
- **Match traffic based on URL category for policy enforcement**—If you want a specific policy rule to apply only to web traffic to sites in a specific category, you would add the category as match criteria when you create the policy rule. For example, you could use the URL category *streaming-media* in a QoS policy to apply bandwidth controls to all websites that are categorized as streaming media. See URL Category as Policy Match Criteria for more information.

By grouping websites into categories, it makes it easy to define actions based on certain types of websites. In addition to the standard URL categories, there are three additional categories:

| | |
|---|---|
| not-resolved | Indicates that the website was not found in the local URL filtering database and the firewall was unable to connect to the cloud database to check the category. When a URL category lookup is performed, the firewall first checks the dataplane cache for the URL; if no match is found, it checks the management plane cache, and if no match is found there, it queries the URL database in the cloud. In the case of the PAN-DB private cloud, the URL database in the cloud is not used for queries.<br>Setting the action to block for traffic that is categorized as *not-resolved*, may be very disruptive to users. You could set the action as **continue**, so that users you can notify users that they are accessing a site that is blocked by company policy and provide the option to read the disclaimer and continue to the website.<br>For more information on troubleshooting lookup issues, see Troubleshoot URL Filtering. |
| private-ip-addresses | Indicates that the website is a single domain (no sub-domains), the IP address is in the private IP range, or the URL root domain is unknown to the cloud. |
| unknown | The website has not yet been categorized, so it does not exist in the URL filtering database on the firewall or in the URL cloud database.<br>When deciding on what action to take for traffic categorized as *unknown*, be aware that setting the action to block may be very disruptive to users because there could be a lot of valid sites that are not in the URL database yet. If you do want a very strict policy, you could block this category, so websites that do not exist in the URL database cannot be accessed.<br>See Configure URL Filtering. |

 © Palo Alto Networks, Inc.

## URL Filtering Profile

A URL filtering profile is a collection of URL filtering controls that are applied to individual security policy rules to enforce your web access policy. The firewall comes with a default profile that is configured to block threat-prone categories, such as malware, phishing, and adult. You can use the default profile in a security policy, clone it to be used as a starting point for new URL filtering profiles, or add a new URL filtering profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories. For example, you may want to block social-networking sites, but allow some websites that are part of the social-networking category.

- URL Filtering Profile Actions
- Block and Allow Lists
- Safe Search Enforcement
- Container Pages
- HTTP Header Logging

© Palo Alto Networks, Inc.

## URL Filtering Profile Actions

The URL Filtering profile specifies an action for each known URL category. By default, all URL categories are set to allow when you Create a new URL Filtering profile. This means that the users will be able to browse to all sites freely and the traffic will not be logged. The firewall also comes predefined default URL filtering profile that allows access to all categories except the following threat-prone categories, which it blocks: abused-drugs, adult, gambling, hacking, malware, phishing, questionable, and weapons.

As a best practice, if you want to create a custom URL Filtering category, clone the default URL filtering profile and change the action in all allow categories to either alert or continue so that you have visibility into the traffic. It is also a best practice to set the proxy-avoidance-and-anonymizers category to block.

| Action | Description |
|---|---|
| **alert** | The website is allowed and a log entry is generated in the URL filtering log. |
| **allow** | The website is allowed and no log entry is generated. |
| **block** | The website is blocked and the user will see a response page and will not be able to continue to the website. A log entry is generated in the URL filtering log. |
| **continue** | The user will be prompted with a response page indicating that the site has been blocked due to company policy, but the user is prompted with the option to continue to the website. The **continue** action is typically used for categories that are considered benign and is used to improve the user experience by giving them the option to continue if they feel the site is incorrectly categorized. The response page message can be customized to contain details specific to your company. A log entry is generated in the URL filtering log.<br> The Continue page will not be displayed properly on client machines that are configured to use a proxy server. |
| **override** | The user will see a response page indicating that a password is required to allow access to websites in the given category. With this option, the security admin or helpdesk person would provide a password granting temporary access to all websites in the given category. A log entry is generated in the URL filtering log. See Configure URL Admin Override.<br> The Override page will not be displayed properly on client machines that are configured to use a proxy server. |
| **none** | The **none** action only applies to custom URL categories. Select **none** to ensure that if multiple URL profiles exist, the custom category will not have any impact on other profiles. For example, if you have two URL profiles and the custom URL category is set to **block** in one of the profiles, the other profile should have the action set to **none** if you do not want it to apply.<br>Also, in order to delete a custom URL category, it must be set to **none** in any profile where it is used. |

 © Palo Alto Networks, Inc.

## Block and Allow Lists

In some cases you might want to block a category, but allow a few specific sites in that category. Alternatively, you might want to allow some categories, but block individual sites in the category. You do this by adding the IP addresses or URLs of these sites in the Block list and Allow list sections of the URL Filtering profile to Define websites that should always be blocked or allowed. When entering URLs in the Block List or Allow List, enter each URL or IP address in a new row separated by a new line. When using wildcards in the URLs, follow these rules:

- Do not include HTTP and HTTPS when defining URLs. For example, enter www.paloaltonetworks.com or paloaltonetworks.com instead of https://www.paloaltonetworks.com.
- Entries in the block list must be an exact match and are case-insensitive.

  For example: If you want to prevent a user from accessing any website within the domain paloaltonetworks.com, you would also add *.paloaltonetworks.com, so whatever domain prefix (http://, www, or a sub-domain prefix such as mail.paloaltonetworks.com) is added to the address, the specified action will be taken. The same applies to the sub-domain suffix; if you want to block paloaltonetworks.com/en/US, you would need to add paloaltonetworks.com/* as well.

  Further, if you want to limit access to a domain suffix such as `paloaltonetworks.com.au`, you must add a `/`, so that the match restricts a dot that follows `.com.` In this case, you need to add the entry as `*.paloaltonetworks.com/`

  Block and allow lists support wildcard patterns. The following characters are considered separators:

  .

  /

  ?

  &

  =

  ;

  +

  Every substring that is separated by the characters listed above is considered a token. A token can be any number of ASCII characters that does not contain any separator character or *. For example, the following patterns are valid:

  *.yahoo.com (tokens are: "*", "yahoo" and "com")

  www.*.com (tokens are: "www", "*" and "com")

  www.yahoo.com/search=* (tokens are: "www", "yahoo", "com", "search", "*")

  The following patterns are invalid because the character "*" is not the only character in the token.

  ww*.yahoo.com

  www.y*.com

## Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos in search query return traffic. On the firewall, you can Enable Safe Search Enforcement so that the firewall will block search results if the end user is not using the strictest safe search settings in the search query. The firewall can enforce safe search for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. This is a best-effort setting and is not guaranteed by the search providers to work with every website.

To use this feature you must enable the **Safe Search Enforcement** option in a URL filtering profile and attach it to a security policy rule.The firewall will then block any matching search query return traffic that is not using the strictest safe search settings. There are two methods for blocking the search results:

- Block Search Results that are not Using Strict Safe Search Settings—When an end user attempts to perform a search without first enabling the strictest safe search settings, the firewall blocks the search query results and displays the URL Filtering Safe Search Block Page. By default, this page will provide a URL to the search provider settings for configuring safe search.
- Enable Transparent Safe Search Enforcement—When an end user attempts to perform a search without first enabling the strict safe search settings, the firewall blocks the search results with an HTTP 503 status code and redirects the search query to a URL that includes the safe search parameters. You enable this functionality by importing a new URL Filtering Safe Search Block Page containing the JavaScript for rewriting the search URL to include the strict safe search parameters. In this configuration, users will not see the block page, but will instead be automatically redirected to a search query that enforces the strictest safe search options. This safe search enforcement method requires content release version 475 or later and is only supported for Google, Yahoo, and Bing searches.

Also, because most search providers now use SSL to return search results, you must also configure a Decryption policy rule for the search traffic to enable the firewall to inspect the search traffic and enforce safe search.

Safe search enforcement enhancements and support for new search providers is periodically added in content releases. This information is detailed in the Application and Threat Content Release Notes. How sites are judged to be safe or unsafe is performed by each search provider, not by Palo Alto Networks.

Safe search settings differ by search provider as detailed in Table: Search Provider Safe Search Settings.

 © Palo Alto Networks, Inc.

**Table: Search Provider Safe Search Settings**

| Search Provider | Safe Search Setting Description |
|---|---|
| Google/YouTube | Offers safe search on individual computers or network-wide through Google's safe search virtual IP address:<br><br>**Safe Search Enforcement for Google Searches on Individual Computers**<br><br>In the Google Search Settings, the **Filter explicit results** setting enables safe search functionality. When enabled, the setting is stored in a browser cookie as `FF=` and passed to the server each time the user performs a Google search.<br><br>Appending `safe=active` to a Google search query URL also enables the strictest safe search settings.<br><br>**Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address**<br><br>Google provides servers that Lock SafeSearch (forcesafesearch.google.com) settings in every Google and YouTube search. By adding a DNS entry for `www.google.com` and `www.youtube.com` (and other relevant Google and YouTube country subdomains) that includes a CNAME record pointing to `forcesafesearch.google.com` to your DNS server configuration, you can ensure that all users on your network are using strict safe search settings every time they perform a Google or YouTube search. Keep in mind, however, that this solution is not compatible with Safe Search Enforcement on the firewall. Therefore, if you are using this option to force safe search on Google, the best practice is to block access to other search engines on the firewall by creating custom URL categories and adding them to the block list in the URL filtering profile.<br><br>If you plan to use the Google Lock SafeSearch solution, consider configuring DNS Proxy (**Network > DNS Proxy**) and setting the inheritance source as the Layer 3 interface on which the firewall receives DNS settings from service provider via DHCP. You would configure the DNS proxy with **Static Entries** for www.google.com and www.youtube.com, using the local IP address for the forcesafesearch.google.com server. |
| Yahoo | Offers safe search on individual computers only. The Yahoo Search Preferences includes three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `vm=` and passed to the server each time the user performs a Yahoo search.<br><br>Appending `vm=r` to a Yahoo search query URL also enables the strictest safe search settings.<br><br>When performing a search on Yahoo Japan (yahoo.co.jp) while logged into a Yahoo account, end users must also enable the **SafeSearch Lock** option. |
| Bing | Offers safe search on individual computers or through their Bing in the Classroom program. The Bing Settings include three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `adlt=` and passed to the server each time the user performs a Bing search.<br><br>Appending `adlt=strict` to a Bing search query URL also enables the strictest safe search settings.<br><br>The Bing SSL search engine does not enforce the safe search URL parameters and you should therefore consider blocking Bing over SSL for full safe search enforcement. |

## Container Pages

A container page is the main page that a user accesses when visiting a website, but additional websites may be loaded within the main page. If the **Log Container page only** option is enabled in the URL filtering profile, only the main container page will be logged, not subsequent pages that may be loaded within the container page. Because URL filtering can potentially generate a lot of log entries, you may want to turn on this option, so log entries will only contain those URIs where the requested page file name matches the specific mime-types. The default set includes the following mime-types:

- application/pdf
- application/soap+xml
- application/xhtml+xml
- text/html
- text/plain
- text/xml

If you have enabled the **Log container page only** option, there may not always be a correlated URL log entry for threats detected by antivirus or vulnerability protection.

 © Palo Alto Networks, Inc.

## HTTP Header Logging

URL filtering provides visibility and control over web traffic on your network. For improved visibility into web content, you can configure the URL Filtering profile to log HTTP header attributes included in a web request. When a client requests a web page, the HTTP header includes the user agent, referer, and x-forwarded-for fields as attribute-value pairs and forwards them to the web server. When enabled for logging HTTP headers, the firewall logs the following attribute-value pairs in the URL Filtering logs:

| Attribute | Description |
|---|---|
| User-Agent | The web browser that the user used to access the URL, for example, Internet Explorer. This information is sent in the HTTP request to the server. |
| Referer | The URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested. |
| X-Forwarded-For (XFF) | The option in the HTTP request header field that preserves the IP address of the user who requested the web page. If you have a proxy server on your network, the XFF allows you to identify the IP address of the user who requested the content, instead of only recording the proxy server's IP address as source IP address that requested the web page. |

## URL Filtering Response Pages

The firewall provides three predefined response pages that display by default when a user attempts to browse to a site in a category that is configured with one of the block actions in the URL Filtering Profile (block, continue, or override) or when Safe Search Enforcement is enabled:

- **URL Filtering and Category Match Block Page**—Access blocked by a URL Filtering Profile or because the URL category is blocked by a security policy.

- **URL Filtering Continue and Override Page**—Page with initial block policy that allows users to bypass the block by clicking **Continue**. With URL Admin Override enabled, (Configure URL Admin Override), after clicking **Continue**, the user must supply a password to override the policy that blocks the URL.

- **URL Filtering Safe Search Block Page**—Access blocked by a security policy with a URL filtering profile that has the Safe Search Enforcement option enabled (see Enable Safe Search Enforcement). The user will see this page if a search is performed using Google, Bing, Yahoo, or Yandex and their browser or search engine account setting for Safe Search is not set to strict.

You can either use the predefined pages, or you can Customize the URL Filtering Response Pages to communicate your specific acceptable use policies and/or corporate branding. In addition, you can use the URL Filtering Response Page Variables for substitution at the time of the block event or add one of the supported Response Page References to external images, sounds, or style sheets.

 © Palo Alto Networks, Inc.

**URL Filtering Response Page Variables**

| Variable | Usage |
|---|---|
| <user/> | The firewall replaces the variable with the username (if available via User-ID) or IP address of the user when displaying the response page. |
| <url/> | The firewall replaces the variable with the requested URL when displaying the response page. |
| <category/> | The firewall replaces the variable with the URL filtering category of the blocked request. |
| <pan_form/> | HTML code for displaying the **Continue** button on the URL Filtering Continue and Override page. |

You can also add code that triggers the firewall to display different messages depending on what URL category the user is attempting to access. For example, the following code snippet from a response page specifies to display Message 1 if the URL category is games, Message 2 if the category is travel, or Message 3 if the category is kids:

```
var cat = "<category/>";
switch(cat)
{
   case 'games':
      document.getElementById("warningText").innerHTML = "Message 1";
      break;
   case 'travel':
      document.getElementById("warningText").innerHTML = "Message 2";
      break;
   case 'kids':
      document.getElementById("warningText").innerHTML = "Message 3";
      break;
}
```

Only a single HTML page can be loaded into each virtual system for each type of block page. However, other resources such as images, sounds, and cascading style sheets (CSS files) can be loaded from other servers at the time the response page is displayed in the browser. All references must include a fully qualified URL.

**Response Page References**

| Reference Type | Example HTML Code |
|---|---|
| Image | `<img src="http://virginiadot.org/images/Stop-Sign-gif.gif">` |
| Sound | `<embed src="http://simplythebest.net/sounds/WAV/WAV_files/ movie_WAV_files/ do_not_go.wav" volume="100" hidden="true" autostart="true">` |
| Style Sheet | `<link href="http://example.com/style.css" rel="stylesheet" type="text/css" />` |
| Hyperlink | `<a href="http://en.wikipedia.org/wiki/Acceptable_use_policy">View Corporate Policy</a>` |

## URL Category as Policy Match Criteria

Use URL Categories as a match criteria in a policy rule for more granular enforcement. For example, suppose you have configured Decryption, but you want to exclude traffic to certain types of websites (for example, healthcare or financial services) from being decrypted. In this case you could create a decryption policy rule that matches those categories and set the action to no-decrypt. By placing this rule above the rule to decrypt all trafic, you can ensure that web traffic with URL categories that match the no-decrypt rule, and all other traffic would match the subsequent rule.

The following table describes the policy types that accept URL category as match criteria:

| Policy Type | Description |
|---|---|
| Captive Portal | To ensure that users authenticate before being allowed access to a specific category, you can attach a URL category as a match criterion for the Captive Portal policy. |
| Decryption | Decryption policies can use URL categories as match criteria to determine if specified websites should be decrypted or not. For example, if you have a decryption policy with the action decrypt for all traffic between two zones, there may be specific website categories, such as *financial-services* and/or *health-and-medicine*, that should not be decrypted. In this case, you would create a new decryption policy with the action of *no-decrypt* that precedes the decrypt policy and then defines a list of URL categories as match criteria for the policy. By doing this, each URL category that is part of the no-decrypt policy will not be decrypted. You could also configure a custom URL category to define your own list of URLs that can then be used in the no-decrypt policy. |
| QoS | QoS policies can use URL categories to allocate throughput levels for specific website categories. For example, you may want to allow the streaming-media category, but limit throughput by adding the URL category as match criteria to the QoS policy. |
| Security | In security policies you can use URL categories both as a match criteria in the **Service/URL Category** tab, and in URL filtering profiles that are attached in the **Actions** tab.<br>If for example, the IT-security group in your company needs access to the hacking category, while all other users are denied access to the category, you must create the following rules:<br>• A security rule that allows the IT-Security group to access content categorized as hacking. The security rule references the *hacking* category in the S**ervices/URL Category** tab and IT-Security group in the **Users** tab.<br>• Another security rule that allows general web access for all users. To this rule you attach a URL filtering profile that blocks the hacking category.<br>The policy that allows access to hacking must be listed before the policy that blocks hacking. This is because security policy rules are evaluated top down, so when a user who is part of the security group attempts to access a hacking site, the policy rule that allows access is evaluated first and will allow the user access to the hacking sites. Users from all other groups are evaluated against the general web access rule which blocks access to the hacking sites. |

 © Palo Alto Networks, Inc.

# PAN-DB Categorization

▲ PAN-DB URL Categorization Components
▲ PAN-DB URL Categorization Workflow

## PAN-DB URL Categorization Components

The following table describes the PAN-DB components in detail. The BrightCloud system works similarly, but does not use an initial seed database.

| Component | Description |
|---|---|
| URL Filtering Seed Database | The initial seed database downloaded to the firewall is a small subset of the database that is maintained on the Palo Alto Networks URL cloud servers. The reason this is done is because the full database contains millions of URLs and many of these URLs may never be accessed by your users. When downloading the initial seed database, you select a region (North America, Europe, APAC, Japan). Each region contains a subset of URLs most accessed for the given region. This allows the firewall to store a much smaller URL database for better URL lookup performance. If a user accesses a website that is not in the local URL database, the firewall queries the full cloud database and then adds the new URL to the local database. This way the local database on the firewall is continually populated/customized based on actual user activity.<br>Note that re-downloading the PAN-DB seed database or switching the URL database vendor from PAN-DB to BrightCloud will clear the local database. |
| Cloud Service<br>See Differences Between the PAN-DB Public Cloud and PAN-DB Private Cloud, for information on the private cloud. | The PAN-DB cloud service is implemented using Amazon Web Services (AWS). AWS provides a distributed, high-performance, and stable environment for seed database downloads and URL lookups for Palo Alto Networks firewalls and communication is performed over SSL. The AWS cloud systems hold the entire PAN-DB and is updated as new URLs are identified. The PAN-DB cloud service supports an automated mechanism to update the firewall's local URL database if the version does not match. Each time the firewall queries the cloud servers for URL lookups, it will also check for critical updates. If there have been no queries to the cloud servers for more than 30 minutes, the firewall will check for updates on the cloud systems.<br>The cloud system also provides a mechanism to submit URL category change requests. This is performed through the test-a-site service and is available directly from the device (URL filtering profile setup) and from the Palo Alto Networks Test A Site website. You can also submit a URL categorization change request directly from the URL filtering log on the firewall in the log details section. |

| Component | Description |
|---|---|
| Management Plane (MP) URL Cache | When you activate PAN-DB on the firewall, the firewall downloads a seed database from one of the PAN-DB cloud servers to initially populate the local cache for improved lookup performance. Each regional seed database contains the top URLs for the region and the size of the seed database (number of URL entries) also depends on the platform. The URL MP cache is automatically written to the firewall's local drive every eight hours, before the firewall is rebooted, or when the cloud upgrades the URL database version on the firewall. After rebooting the firewall, the file that was saved to the local drive will be loaded to the MP cache. A least recently used (LRU) mechanism is also implemented in the URL MP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs. |
| Dataplane (DP) URL Cache | This is a subset of the MP cache and is a customized, dynamic URL database that is stored in the dataplane (DP) and is used to improve URL lookup performance. The URL DP cache is cleared at each firewall reboot. The number of URLs that are stored in the URL DP cache varies by hardware platform and the current URLs stored in the TRIE (data structure). A least recently used (LRU) mechanism is implemented in the DP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs. Entries in the URL DP cache expire after a specified period of time and the expiration period cannot be changed by the administrator. |

 © Palo Alto Networks, Inc.

## PAN-DB URL Categorization Workflow

When a user attempts to access a URL and the URL category needs to be determined, the firewall will compare the URL with the following components (in order) until a match has been found:

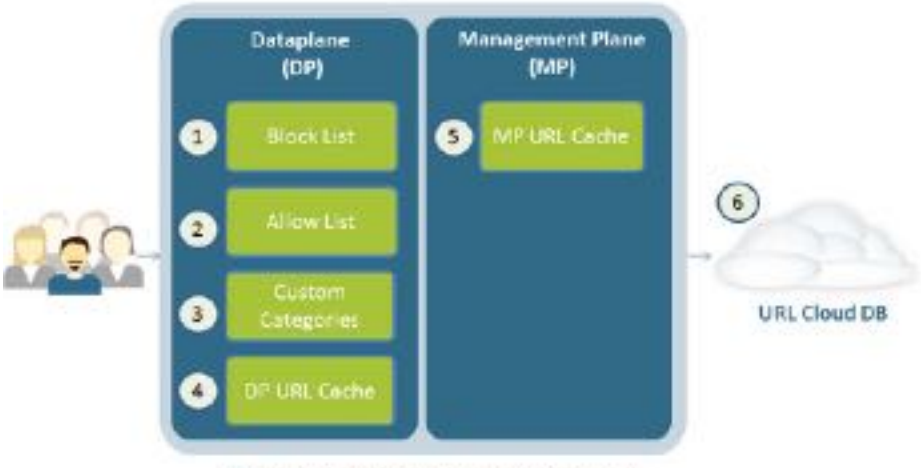

URL Categorization Resolution Process Flow

If a URL query matches an expired entry in the URL DP cache, the cache responds with the expired category, but also sends a URL categorization query to the management plane. This is done to avoid unnecessary delays in the DP, assuming that the frequency of changing categories is low. Similarly, in the URL MP cache, if a URL query from the DP matches an expired entry in the MP, the MP responds to the DP with the expired category and will also send a URL categorization request to the cloud service. Upon getting the response from the cloud, the firewall will resend the updated response to the DP.

As new URLs and categories are defined or if critical updates are needed, the cloud database will be updated. Each time the firewall queries the cloud for a URL lookup or if no cloud lookups have occurred for 30 minutes, the database versions on the firewall be compared and if they do not match, an incremental update will be performed.

© Palo Alto Networks, Inc.

# Enable a URL Filtering Vendor

To enable URL filtering on a firewall, you must purchase and activate a URL Filtering license for one of the supported URL Filtering Vendors and then install the database for the vendor you selected.

Starting with PAN-OS 6.0, firewalls managed by Panorama do not need to be running the same URL filtering vendor that is configured on Panorama. For firewalls running PAN-OS 6.0 or later, when a mismatch is detected between the vendor enabled on the firewalls and what is enabled on Panorama, the firewalls can automatically migrate URL categories and/or URL profiles to (one or more) categories that align with that of the vendor enabled on it. For guidance on how to configure URL Filtering on Panorama if you are managing firewalls running different PAN-OS versions, refer to the Panorama Administrator's Guide.

If you have valid licenses for both PAN-DB and BrightCloud, activating the PAN-DB license automatically deactivates the BrightCloud license (and vice versa). At a time, only one URL filtering license can be active on a firewall.

- Enable PAN-DB URL Filtering
- Enable BrightCloud URL Filtering


© Palo Alto Networks, Inc.

# Enable PAN-DB URL Filtering

| Enable PAN-DB URL Filtering | |
|---|---|
| Step 1 Obtain and install a PAN-DB URL filtering license and confirm that it is installed.<br><br>If the license expires, PAN-DB URL Filtering continues to work based on the URL category information that exists in the dataplane and management plane caches. However, URL cloud lookups and other cloud-based updates will not function until you install a valid license. | 1. Select **Device > Licenses** and, in the License Management section, select the license installation method:<br>• **Retrieve license keys from license server**<br>• **Activate feature using authorization code**<br>• **Manually upload license key**<br>2. After installing the license, confirm that the PAN-DB URL Filtering section, **Date Expires** field, displays a valid date. |
| Step 2 Download the initial seed database and activate PAN-DB URL Filtering.<br><br>The firewall must have Internet access; you cannot manually upload the PAN-DB seed database. | 1. In the PAN-DB URL Filtering section, **Download Status** field, click **Download Now**.<br>2. Choose a region (North America, Europe, APAC, Japan) and then click **OK** to start the download.<br>3. After the download completes, click **Activate**.<br><br>If PAN-DB is already the active URL filtering vendor and you click **Re-Download**, this will reactivate PAN-DB by clearing the dataplane and management plane caches and replacing them with the contents of the new seed database. You should avoid doing this unless it is necessary, as you will lose your cache, which is customized based on the web traffic that has previously passed through the firewall based on user activity. |
| Step 3 Schedule the firewall to download dynamic updates for Applications and Threats.<br><br>A Threat Prevention license is required to receive content updates, which covers Antivirus and Applications and Threats. | 1. Select **Device > Dynamic Updates**.<br>2. In the Schedule field in the Applications and Threats section, click the **None** link to schedule periodic updates.<br><br>You can only schedule dynamic updates if the firewall has direct Internet access. If updates are already scheduled in a section, the link text displays the schedule settings.<br><br>The Applications and Threats updates might contain updates for URL filtering related to the **Safe Search Enforcement** option in the URL filtering profile (**Objects > Security Profiles > URL Filtering**). For example, if Palo Alto Networks adds support for a new search provider vendor or if the method used to detect the Safe Search setting for an existing vendor changes, the Application and Threats updates will include that update. |

# Enable BrightCloud URL Filtering

| Enable BrightCloud URL Filtering | | |
|---|---|---|
| Step 1 | Obtain and install a BrightCloud URL filtering license and confirm that it is installed.<br> BrightCloud has an option in the URL filtering profile (**Objects > Security Profiles > URL Filtering**) to either allow all categories or block all categories if the license expires. | 1. Select **Device > Licenses** and, in the **License Management** section, select the license installation method:<br>• **Activate feature using authorization code**<br>• **Retrieve license keys from license server**<br>• **Manually upload license key**<br>2. After installing the license, confirm that the BrightCloud URL Filtering section, **Date Expires** field, displays a valid date. |
| Step 2 | Install the BrightCloud database.<br>The way you do this depends on whether or not the firewall has direct Internet access. | **Firewall with Direct Internet Access**<br>Select **Device > Licenses** and in the BrightCloud URL Filtering section, Active field, click the **Activate** link to install the BrightCloud database. This operation automatically initiates a system reset.<br>**Firewall without Direct Internet Access**<br>1. Download the BrightCloud database to a host that has Internet access. The firewall must have access to the host:<br>a. On a host with Internet access, go to the Palo Alto Support website (https://support.paloaltonetworks.com) and log in.<br>b. In the Resources section, click **Dynamic Updates**.<br>c. In the BrightCloud Database section, click **Download** and save the file to the host.<br>2. Upload the database to the firewall:<br>a. Log in to the firewall, select **Device > Dynamic Updates** and click **Upload**.<br>b. For the **Type**, select **URL Filtering**.<br>c. Enter the path to the **File** on the host or click **Browse** to find it, then click **OK**. When the Status is **Completed**, click **Close**.<br>3. Install the database:<br>a. Select **Device > Dynamic Updates** and click **Install From File**.<br>b. For the **Type**, select **URL Filtering**. The firewall automatically selects the file you just uploaded.<br>c. Click **OK** and, when the Result is **Succeeded**, click **Close**. |
| Step 3 | Enable cloud lookups for dynamically categorizing a URL if the category is not available on the local BrightCloud database. | 1. Access the PAN-OS CLI.<br>2. Enter the following commands to enable dynamic URL filtering:<br>`configure`<br>`set deviceconfig setting url dynamic-url yes`<br>`commit` |

 © Palo Alto Networks, Inc.

| Enable BrightCloud URL Filtering (Continued) | |
|---|---|
| Step 4 Schedule the firewall to download dynamic updates for Applications and Threats signatures and URL filtering.<br><br>You can only schedule dynamic updates if the firewall has direct Internet access.<br><br>The Applications and Threats updates might contain updates for URL filtering related to the **Safe Search Enforcement** option in the URL filtering profile. For example, if Palo Alto Networks adds support for a new search provider vendor or if the method used to detect the Safe Search setting for an existing vendor changes, the Application and Threats updates will include that update.<br><br>BrightCloud updates include a database of approximately 20 million websites that are stored locally on the firewall. You must schedule URL filtering updates to receive BrightCloud database updates.<br><br> A Threat Prevention license is required to receive Antivirus and Applications and Threats updates. | 1. Select **Device > Dynamic Updates**.<br>2. In the Applications and Threats section, Schedule field, click the **None** link to schedule periodic updates.<br>3. In the URL Filtering section, Schedule field, click the **None** link to schedule periodic updates.<br><br> If updates are already scheduled in a section, the link text displays the schedule settings. |

# Determine URL Filtering Policy Requirements

The recommended practice for deploying URL filtering in your organization is to first start with a passive URL filtering profile that will alert on most categories. After setting the alert action, you can then monitor user web activity for a few days to determine patterns in web traffic. After doing so, you can then make decisions on the websites and website categories that should be controlled.

In the procedure that follows, threat-prone sites will be set to block and the other categories will be set to alert, which will cause all websites traffic to be logged. This may potentially create a large amount of log files, so it is best to do this for initial monitoring purposes to determine the types of websites your users are accessing. After determining the categories that your company approves of, those categories should then be set to allow, which will not generate logs. You can also reduce URL filtering logs by enabling the **Log container page only** option in the URL Filtering profile, so only the main page that matches the category will be logged, not subsequent pages/categories that may be loaded within the container page.

| **Configure and Apply a Passive URL Filtering Profile** | | |
|---|---|---|
| Step 1 | Create a new URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select the default profile and then click **Clone**. The new profile will be named **default-1**.<br>3. Select the **default-1** profile and rename it. For example, rename it to URL-Monitoring. |
| Step 2 | Configure the action for all categories to **alert**, except for threat-prone categories, which should remain blocked.<br><br>To select all items in the category list from a Windows system, click the first category, then hold down the shift key and click the last category—this will select all categories. Hold the control key (ctrl) down and click items that should be deselected. On a Mac, do the same using the shift and command keys. You could also just set all categories to alert and manually change the recommended categories back to block. | 1. In the section that lists all URL categories, select all categories.<br>2. To the right of the *Action* column heading, mouse over and select the down arrow and then select **Set Selected Actions** and choose **alert**.<br> 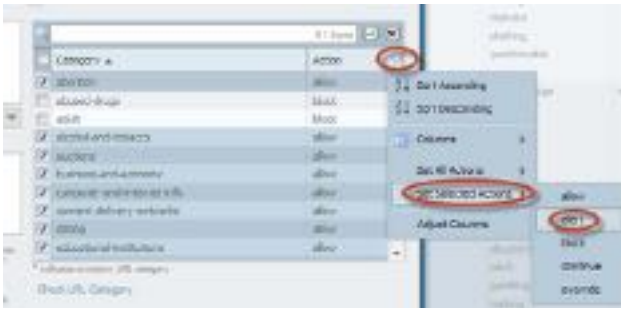 <br>3. To ensure that you block access to threat-prone sites, select the following categories and then set the action to **block**: abused-drugs, adult, gambling, hacking, malware. phishing, questionable, weapons.<br>4. Click **OK** to save the profile. |
| Step 3 | Apply the URL Filtering profile to the security policy rule(s) that allows web traffic for users. | 1. Select **Policies > Security** and select the appropriate security policy to modify it.<br>2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the new profile.<br>3. Click **OK** to save. |
| Step 4 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

| Configure and Apply a Passive URL Filtering Profile (Continued) | | |
|---|---|---|
| Step 5 | View the URL filtering logs to determine all of the website categories that your users are accessing. In this example, some categories are set to block, so those categories will also appear in the logs.<br><br>For information on viewing the logs and generating reports, see Monitor Web Activity. | Select **Monitor > Logs > URL Filtering**. A log entry will be created for any website that exists in the URL filtering database that is in a category that is set to any action other than **allow**. |

# Monitor Web Activity

The ACC, URL filtering logs and reports show all user web activity for URL categories that are set to **alert**, **block**, **continue**, or **override**. By monitoring the logs, you can gain a better understanding of the web activity of your user base to determine a web access policy.

The following topics describe how to monitor web activity:

- Monitor Web Activity of Network Users
- View the User Activity Report
- Configure Custom URL Filtering Reports


© Palo Alto Networks, Inc.

## Monitor Web Activity of Network Users

You can use the ACC, and the URL filtering reports and logs that are generated on the firewall to track user activity.

For a quick view of the most common categories users access in your environment, check the **ACC** widgets. Most widgets in the Network Activity tab, allows you to sort on URLs. For example, in the Application Usage widget, you can see that the networking category is the most accessed category, followed by encrypted tunnel, and ssl. You can also view the list of **Threat Activity** and **Blocked Activity** sorted on URLs.

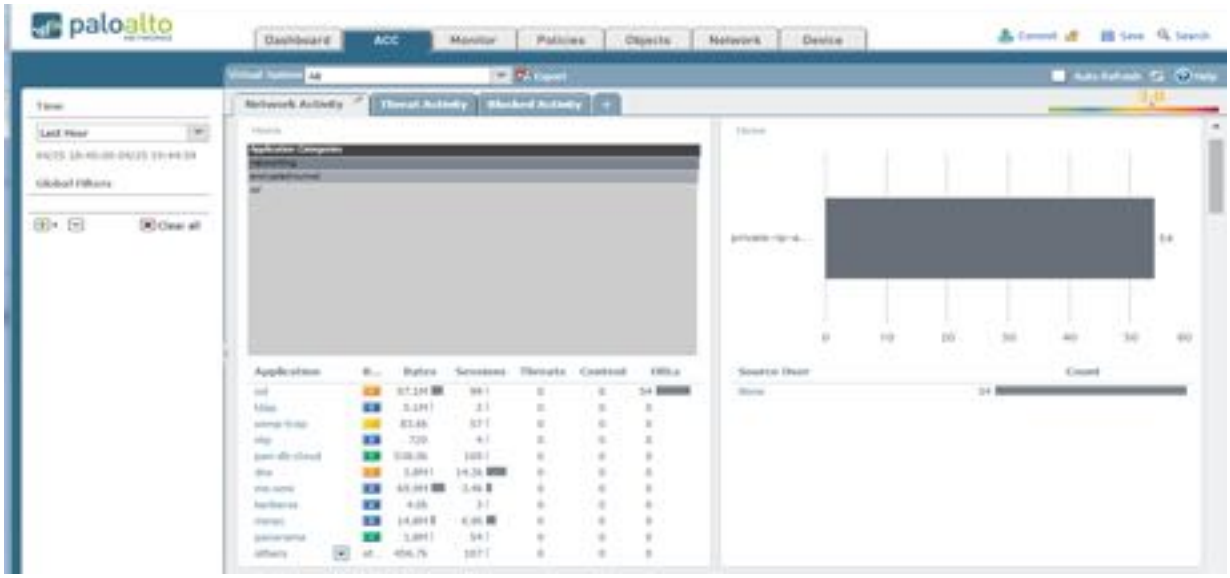

From the ACC, you can directly **Jump to the Logs** or you can navigate to **Monitor > Logs > URL filtering** to view the URL filtering logs. The following bullet points show examples of the URL filtering logs ().

- **Alert log**—In this log, the category is shopping and the action is alert.

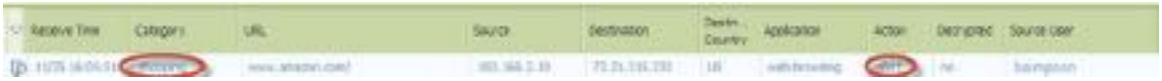

- **Block log**—In this log, the category alcohol-and-tobacco was set to block, so the action is block-url and the user will see a response page indicating that the website was blocked.

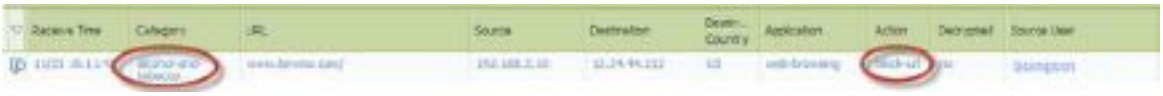

- **Alert log on encrypted website**—In this example, the category is social-networking and the application is facebook-base, which is required to access the Facebook website and other Facebook applications. Because faceboook.com is always encrypted using SSL, the traffic was decrypted by the firewall, which allows the website to be recognized and controlled if needed.

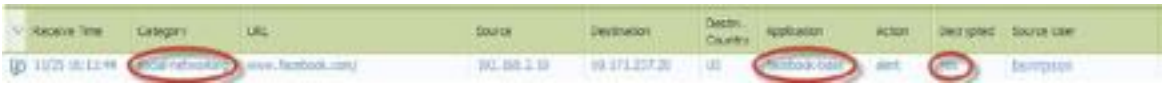

You can also add several other columns to your URL Filtering log view, such as: to and from zone, content type, and whether or not a packet capture was performed. To modify what columns to display, click the down arrow in any column and select the attribute to display.

To view the complete log details and/or request a category change for the given URL that was accessed, click the log details icon in the first column of the log.

To generate a predefined URL filtering reports on URL categories, URL users, Websites accessed, Blocked categories, and more, select **Monitor > Reports** and under the **URL Filtering Reports** section, select one of the reports. The reports are based on a 24-hour period and the day is selected by choosing a day in the calendar section. You can also export the report to PDF, CSV, or XML.

 © Palo Alto Networks, Inc.

Report shows computer-and-internet-info is the most popular category
Select a report
Select a date

## View the User Activity Report

This report provides a quick method of viewing user or group activity and also provides an option to view browse time activity.

| Generate a User Activity Report | |
|---|---|
| Step 1 Configure a User Activity Report | 1. Select **Monitor > PDF Reports > User Activity Report**.<br>2. Enter a report **Name** and select the report type. Select **User** to generate a report for one person, or select **Group** for a group of users.<br>You must Enable User-ID in order to be able to select user or group names. If User-ID is not configured, you can select the type **User** and enter the IP address of the user's computer.<br>3. Enter the Username/IP address for a user report or enter the group name for a user group report.<br>4. Select the time period. You can select an existing time period, or select **Custom**.<br>5. Select the **Include Detailed Browsing** check box, so browsing information is included in the report.<br> |
| Step 2 Run the user activity report and then download the report. | 1. Click **Run Now**.<br>2. After the report is generated, click the **Download User Activity Report** link.<br> 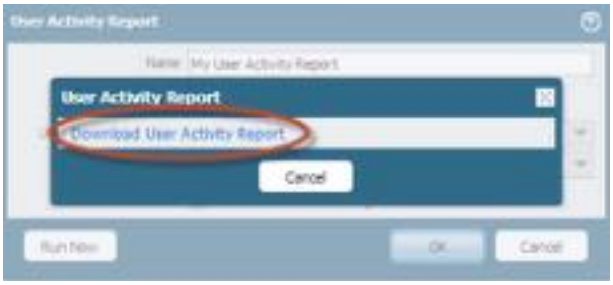 <br>3. After the report is downloaded, click **Cancel** and then click **OK** to save the report. |

 © Palo Alto Networks, Inc.

**Generate a User Activity Report (Continued)**

Step 3 View the user activity report by opening the PDF file that was downloaded. The top of the report will contain a table of contents similar to the following:

User Activity Report for 192.168.2.10
Tuesday, December 31, 2013 09:35:47 - 10:35:46


Application Usage 2
Traffic Summary by URL Category 3
Browsing Summary by URL Category 4
Browsing Summary by Website 5
Blocked Browsing Summary by Website 6
Detailed Web Browsing Activity 7


Step 4 Click an item in the table of contents to view details. For example, click *Traffic Summary by URL Category* to view statistics for the selected user or group.

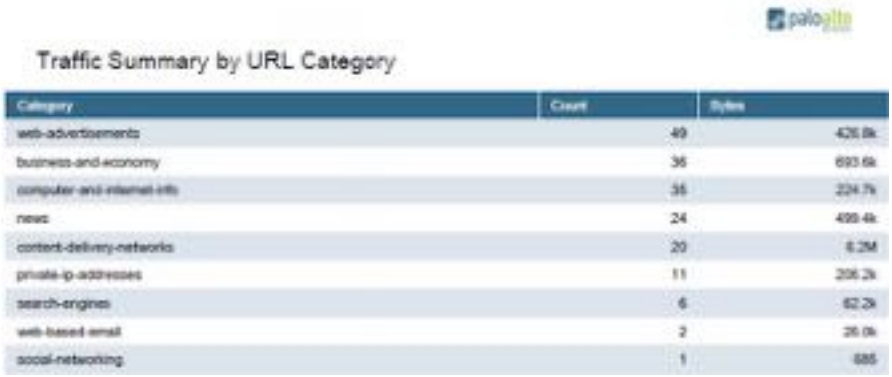
Traffic Summary by URL Category

| Category | Count | Bytes |
|---|---|---|
| web-advertisements | 49 | 426.8k |
| business-and-economy | 36 | 693.6k |
| computer-and-internet-info | 35 | 224.7k |
| news | 24 | 499.4k |
| content-delivery-networks | 20 | 6.2M |
| private-ip-addresses | 11 | 206.2k |
| search-engines | 6 | 62.2k |
| web-based-email | 2 | 26.0k |
| social-networking | 1 | 686 |

## Configure Custom URL Filtering Reports

To generate a detailed report that can also be scheduled, you can configure a custom report and select from a list of all available URL filtering log fields.

| Configure a Custom URL Filtering Report | | |
|---|---|---|
| Step 1 | Add a new custom report. | 1. Select **Monitor > Manage Custom Reports** and click **Add**.<br>2. Enter a report **Name**, for example, My-URL-Custom-Report.<br>3. From the **Database** drop-down, select **URL Log**. |
| Step 2 | Configure report options. | 1. Select the **Time Frame** drop-down and select a range.<br>2. (Optional) To customize how the report is sorted and grouped, select **Sort By** and chose the number of items to display (top 25 for example) and then select **Group By** and select an option such as **Category**, and then select how many groups will be defined.<br>3. In the **Available Columns** list, select the fields to include the report. The following columns are typically used for a URL report:<br>• Action<br>• Category<br>• Destination Country<br>• Source User<br>• URL<br> |
| Step 3 | Run the report to check the results. If the results are satisfactory, set a schedule to run the report automatically. | 1. Click the **Run Now** icon to immediately generate the report that will appear in a new tab.<br>2. (Optional) Click the **Schedule** check box to run the report once per day. This will generate a daily report that details web activity over the last 24 hours. To access the report, select **Monitor > Report** and then expand **Custom Reports** on the right column and select the report. |
| Step 4 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

# Configure URL Filtering

After you Determine URL Filtering Policy Requirements, you should have a basic understanding of what types of websites and website categories your users are accessing. With this information, you are now ready to create custom URL filtering profiles and attach them to the security policy rule(s) that allow web access.

| Configure Website Controls | | |
|---|---|---|
| Step 1 | Create a URL Filtering profile or select an existing one.<br> Because the default URL filtering profile blocks risky and threat-prone content, it is a best practice to clone this profile rather than creating a new profile to preserve these default settings. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>Select the default profile and then click **Clone**. The new profile will be named default-1.<br>2. Select the new profile and rename it. |
| Step 2 | Define how to control access to web content. | In the **Categories** tab, for each category that you want visibility into or control over, select a value from the **Action** column as follows:<br>• If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **allow**.<br>• For visibility into traffic to sites in a category, select **alert**.<br>• To deny access to traffic that matches the category and to enable logging of the blocked traffic, select **block**.<br>• To require users to click Continue to proceed to a questionable site, select **continue**.<br>• To only allow access if users provide a configured password, select override. For more details on this setting, see Configure URL Admin Override. |
| Step 3 | Define websites that should always be blocked or allowed.<br>For example, to reduce URL filtering logs, you may want add you corporate websites in the allow list, so no logs will be generated for those sites. Or, if there is a website this is being overly used and is not work related in any way, you can add it to the block list.<br>Items in the block list will always be blocked regardless of the action for the associated category, and URLs in the allow list will always be allowed.<br>For more information on the proper format and wildcards usage, see Block and Allow Lists. | 1. In the URL filtering profile, enter URLs or IP addresses in the **Block List** and select an action:<br>• **block**—Block the URL.<br>• **continue**—Prompt users click **Continue** to proceed to the web page.<br>• **override**—The user will be a prompted for a password to continue to the website.<br>• **alert**—Allow the user to access the website and add an alert log entry in the URL log.<br>2. For the **Allow** list, enter IP addresses or URLs that should always be allowed. Each row must be separated by a new line.<br>3. (Optional) Enable Safe Search Enforcement. |

| Configure Website Controls | | |
|---|---|---|
| Step 4 | Modify the setting to log Container Pages only. | The **Log container page only** option is enabled by default so that only the main page that matches the category is logged, not subsequent pages/categories that may be loaded within the container page. To enable logging for all pages/categories, clear the **Log container page only** check box. |
| Step 5 | Enable HTTP Header Logging for one or more of the supported HTTP header fields. | To log an HTTP header field, select one or more of the following fields to log:<br>• **User-Agent**<br>• **Referer**<br>• **X-Forwarded-For** |
| Step 6 | Save the URL filtering profile. | 1. Click **OK**.<br> Optionally, you can Customize the URL Filtering Response Pages.<br>2. Click **Commit**.<br> To test the URL filtering configuration, simply access a website in a category that is set to block or continue to see if the appropriate action is performed. |

 © Palo Alto Networks, Inc.

# Customize the URL Filtering Response Pages

The firewall provides three predefined URL Filtering Response Pages that display by default when a user attempts to browse to a site in a category that is configured with one of the block actions in the URL Filtering Profile (block, continue, or override) or when Safe Search Enforcement blocks a search attempt. However, you can create your own custom response pages with your corporate branding, acceptable use policies, links to your internal resources as follows:

| **Customize the URL Filtering Response Pages** | | |
|---|---|---|
| Step 1 | Export the default response page(s). | 1. Select **Device > Response Pages**.<br>2. Select the link for the URL filtering response page you want to modify.<br>3. Click the response page (predefined or shared) and then click the **Export** link and save the file to your desktop. |
| Step 2 | Edit the exported page. | 1. Using the HTML text editor of your choice, edit the page:<br>• If you want the response page to display custom information about the specific user, URL, or category that was blocked, add one or more of the supported URL Filtering Response Page Variables.<br>• If you want to include custom images (such as your corporate logo), a sound, or style sheet, or link to another URL, for example to a document detailing your acceptable web use policy, include one or more of the supported Response Page References.<br>2. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding. For example, in Notepad you would select **UTF-8** from the **Encoding** drop-down in the Save As dialog. |
| Step 3 | Import the customized response page. | 1. Select **Device > Response Pages**.<br>2. Select the link that corresponds to the URL Filtering response page you edited.<br>3. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.<br>4. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select **shared** to make it available to all virtual systems.<br>5. Click **OK** to import the file. |
| Step 4 | Save the new response page(s). | **Commit** the changes. |
| Step 5 | Verify that the new response page displays. | From a browser, go to the URL that will trigger the response page. For example, to see a modified URL Filtering and Category Match response page, browse to URL that your URL filtering policy is set to block. |

# Configure URL Admin Override

In some cases there may be URL categories that you want to block, but allow certain individuals to browse to on occasion. In this case, you would set the category action to **override** and define a URL admin override password in the firewall Content-ID configuration. When users attempt to browse to the category, they will be required to provide the override password before they are allowed access to the site. Use the following procedure to configure URL admin override:

| **Configure URL Admin Override** | |
|---|---|
| Step 1 Set the URL admin override password. | 1. Select **Device > Setup > Content ID**.<br>2. In the **URL Admin Override** section, click **Add**.<br>3. In the **Location** field, select the virtual system to which this password applies.<br>4. Enter the **Password** and **Confirm Password**.<br>5. Select an **SSL/TLS Service Profile**. The profile specifies the certificate that the firewall presents to the user if the site with the override is an HTTPS site. For details, see Configure an SSL/TLS Service Profile.<br>6. Select the **Mode** for prompting the user for the password:<br>• **Transparent**—The firewall intercepts the browser traffic destined for site in a URL category you have set to override and impersonates the original destination URL, issuing an HTTP 401 to prompt for the password. Note that the client browser will display certificate errors if it does not trust the certificate.<br>• **Redirect**—The firewall intercepts HTTP or HTTPS traffic to a URL category set to override and redirects the request to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to prompt for the override password. If you select this option, you must provide the **Address** (IP address or DNS hostname) to which to redirect the traffic.<br>7. Click **OK**. |
| Step 2 (Optional) Set a custom override period. | 1. Edit the URL Filtering section.<br>2. To change the amount of time users can browse to a site in a category for which they have successfully entered the override password, enter a new value in the **URL Admin Override Timeout** field. By default, users can access sites within the category for 15 minutes without re-entering the password.<br>3. To change the amount of time users are blocked from accessing a site set to override after three failed attempts to enter the override password, enter a new value in the **URL Admin Lockout Timeout** field. By default, users are blocked for 30 minutes.<br>4. Click **OK**. |


© Palo Alto Networks, Inc.

| Configure URL Admin Override (Continued) | | |
|---|---|---|
| Step 3 | (Redirect mode only) Create a Layer 3 interface to which to redirect web requests to sites in a category configured for override. | 1. Create a management profile to enable the interface to display the URL Filtering Continue and Override Page response page:<br>a. Select **Network > Interface Mgmt** and click **Add**.<br>b. Enter a **Name** for the profile, select **Response Pages**, and then click **OK**.<br>2. Create the Layer 3 interface. Be sure to attach the management profile you just created (on the **Advanced > Other Info** tab of the Ethernet Interface dialog). |
| Step 4 | (Redirect mode only) To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface to which you are redirecting web requests to a site in a URL category configured for override.You can either generate a self-signed certificate or import a certificate that is signed by an external CA. | To use a self-signed certificate, you must first create a root CA certificate and then use that CA to sign the certificate you will use for URL admin override as follows:<br>1. To create a root CA certificate, select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**. Enter a **Certificate Name**, such as RootCA. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed). Make sure you select the **Certificate Authority** check box and then click **Generate** the certificate.<br>2. To create the certificate to use for URL admin override, click **Generate**. Enter a **Certificate Name** and enter the DNS hostname or IP address of the interface as the **Common Name**. In the **Signed By** field, select the CA you created in the previous step. Add an IP address attribute and specify the IP address of the Layer 3 interface to which you will be redirecting web requests to URL categories that have the override action.<br>3. **Generate** the certificate.<br>4. To configure clients to trust the certificate, select the CA certificate on the **Device Certificates** tab and click **Export**. You must then import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO). |
| Step 5 | Specify which URL categories require an override password to enable access. | 1. Select **Objects > URL Filtering** and either select an existing URL filtering profile or **Add** a new one.<br>2. On the **Categories** tab, set the Action to **override** for each category that requires a password.<br>3. Complete any remaining sections on the URL filtering profile and then click **OK** to save the profile. |
| Step 6 | Apply the URL Filtering profile to the security policy rule(s) that allows access to the sites requiring password override for access. | 1. Select **Policies > Security** and select the appropriate security policy to modify it.<br>2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the profile.<br>3. Click **OK** to save. |
| Step 7 | Save the configuration. | Click **Commit**. |

# Enable Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos for search query return traffic. You can configure Safe Search Enforcement the Palo Alto Networks next-generation firewall to prevent search requests that do not have the strictest safe search settings enabled.

The Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address is not compatible with Safe Search Enforcement on the firewall.

There are two ways to enforce Safe Search on the firewall:

- Block Search Results that are not Using Strict Safe Search Settings
- Enable Transparent Safe Search Enforcement

## Block Search Results that are not Using Strict Safe Search Settings

By default, when you enable safe search enforcement, when a user attempts to perform a search without using the strictest safe search settings, the firewall will block the search query results and display the URL Filtering Safe Search Block Page. This page provides a link to the search settings page for the corresponding search provider so that the end user can enable the safe search settings. If you plan to use this default method for enforcing safe search, you should communicate the policy to your end users prior to deploying the policy. See Table: Search Provider Safe Search Settings for details on how each search provider implements safe search. The default URL Filtering Safe Search Block Page provides a link to the search settings for the corresponding search provider. You can optionally Customize the URL Filtering Response Pages.

Alternatively, to enable safe search enforcement so that it is transparent to your end users, configure the firewall to Enable Transparent Safe Search Enforcement.

 © Palo Alto Networks, Inc.

| Enable Safe Search Enforcement | | |
|---|---|---|
| Step 1 | Enable Safe Search Enforcement in the URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select an existing profile to modify, or clone the default profile to create a new profile.<br>3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.<br>4. (Optional) Restrict users to specific search engines:<br>a. On the **Categories** tab, set the **search-engines** category to **block**.<br>b. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:<br>`www.google.com`<br>`www.bing.com`<br>5. Configure other settings as necessary to:<br>• Define how to control access to web content.<br>• Define websites that should always be blocked or allowed.<br>6. Click **OK** to save the profile. |
| Step 2 | Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet. | 1. Select **Policies > Security** and select a rule to which to apply the URL filtering profile that you just enabled for Safe Search Enforcement.<br>2. On the **Actions** tab, select the **URL Filtering** profile.<br>3. Click **OK** to save the security policy rule. |
| Step 3 | Enable SSL Forward Proxy decryption.<br>Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings. | 1. Add a custom URL category for the search sites:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as SearchEngineDecryption.<br>c. **Add** the following to the Sites list:<br>`www.bing.*`<br>`www.google.*`<br>`search.yahoo.*`<br>d. Click **OK** to save the custom URL category object.<br>2. Follow the steps to Configure SSL Forward Proxy.<br>3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**. |

| Enable Safe Search Enforcement (Continued) | | |
|---|---|---|
| Step 4 | (Optional, but recommended) Block Bing search traffic running over SSL.<br><br>Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL. | 1. Add a custom URL category for Bing:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as EnableBingSafeSearch.<br>c. **Add** the following to the Sites list:<br>`www.bing.com/images/*`<br>`www.bing.com/videos/*`<br>d. Click **OK** to save the custom URL category object.<br>2. Create another URL filtering profile to block the custom category you just created:<br>a. Select **Objects > Security Profiles > URL Filtering**.<br>b. **Add** a new profile and give it a descriptive **Name**.<br>c. Locate the custom category in the Category list and set it to **block**.<br>d. Click **OK** to save the URL filtering profile.<br>3. Add a security policy rule to block Bing SSL traffic:<br>a. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.<br>b. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.<br>c. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.<br>d. Select **TCP** as the **Protocol** and set the **Destination Port** to `443`.<br>e. Click **OK** to save the rule.<br>f. Use the **Move** options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled. |
| Step 5 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

| Enable Safe Search Enforcement (Continued) | |
|---|---|
| Step 6 Verify the Safe Search Enforcement configuration.<br><br>This verification step only works if you are using block pages to enforce safe search. If you are using transparent safe search enforcement, the firewall block page will invoke a URL rewrite with the safe search parameters in the query string. | 1. From a computer that is behind the firewall, disable the strict search settings for one of the supported search providers. For example, on bing.com, click the **Preferences** icon on the Bing menu bar.<br><br> 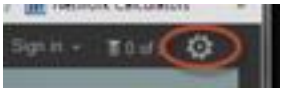 <br><br>2. Set the **SafeSearch** option to **Moderate** or **Off** and click **Save**.<br>3. Perform a Bing search and verify that the URL Filtering Safe Search Block page displays instead of the search results:<br><br> <br><br><br>4. Use the link in the block page to go to the search settings for the search provider and set the safe search setting back to the strictest setting (**Strict** in the case of Bing) and then click **Save**.<br>5. Perform a search again from Bing and verify that the filtered search results display instead of the block page. |

## Enable Transparent Safe Search Enforcement

If you want to enforce filtering of search query results with the strictest safe search filters, but you don't want your end users to have to manually configure the settings, you can enable transparent safe search enforcement as follows. This functionality is supported on Google, Yahoo, and Bing search engines only and requires Content Release version 475 or later.

| **Enable Transparent Safe Search Enforcement** | | |
|---|---|---|
| Step 1 | Make sure the firewall is running Content Release version 475 or later. | 1. Select **Device > Dynamic Updates**.<br>2. Check the **Applications and Threats** section to determine what update is currently running.<br>3. If the firewall is not running the required update or later, click **Check Now** to retrieve a list of available updates.<br>4. Locate the required update and click **Download**.<br>5. After the download completes, click **Install**. |
| Step 1 | Enable Safe Search Enforcement in the URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select an existing profile to modify, or clone the default profile to create a new one.<br>3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.<br>4. (Optional) Allow access to specific search engines only:<br>a. On the **Categories** tab, set the **search-engines** category to **block**.<br>b. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:<br>`www.google.com`<br>`www.bing.com`<br>5. Configure other settings as necessary to:<br>• Define how to control access to web content.<br>• Define websites that should always be blocked or allowed.<br>6. Click **OK** to save the profile. |
| Step 2 | Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet. | 1. Select **Policies > Security** and select a rule to which to apply the URL filtering profile that you just enabled for Safe Search Enforcement.<br>2. On the **Actions** tab, select the **URL Filtering** profile.<br>3. Click **OK** to save the security policy rule. |


© Palo Alto Networks, Inc.

| Enable Transparent Safe Search Enforcement (Continued) | | |
|---|---|---|
| Step 3 | (Optional, but recommended) Block Bing search traffic running over SSL.<br><br>Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL. | 1. Add a custom URL category for Bing:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as EnableBingSafeSearch.<br>c. **Add** the following to the Sites list:<br>`www.bing.com/images/*`<br>`www.bing.com/videos/*`<br>d. Click **OK** to save the custom URL category object.<br>2. Create another URL filtering profile to block the custom category you just created:<br>a. Select **Objects > Security Profiles > URL Filtering**.<br>b. **Add** a new profile and give it a descriptive **Name**.<br>c. Locate the custom category you just created in the Category list and set it to **block**.<br>d. Click **OK** to save the URL filtering profile.<br>3. **Add** a security policy rule to block Bing SSL traffic:<br>a. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.<br>b. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.<br>c. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.<br>d. Select **TCP** as the **Protocol**, set the **Destination Port** to `443`.<br>e. Click **OK** to save the rule.<br>f. Use the **Move** options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled. |
| Step 4 | Edit the URL Filtering Safe Search Block Page, replacing the existing code with the JavaScript for rewriting search query URLs to enforce safe search transparently.<br><br> For a script that you can copy and paste, go here. | 1. Select **Device > Response Pages > URL Filtering Safe Search Block Page**.<br>2. Select **Predefined** and then click **Export** to save the file locally.<br>3. Use an HTML editor and replace all of the existing block page text with the following text and then save the file: |

**Enable Transparent Safe Search Enforcement (Continued)**

```
<html>
<head>
<title>Search Blocked</title>
<meta http-equiv="Content-Type" content="text/html; charset=utf-8">
<META HTTP-EQUIV="PRAGMA" CONTENT="NO-CACHE">
<meta name="viewport" content="initial-scale=1.0">
<style>
  #content {
    border:3px solid#aaa;
    background-color:#fff;
    margin:1.5em;
    padding:1.5em;
    font-family:Tahoma,Helvetica,Arial,sans-serif;
    font-size:1em;
  }
  h1 {
    font-size:1.3em;
    font-weight:bold;
    color:#196390;
  }
  b {
    font-weight:normal;
    color:#196390;
  }
</style>
</head>
    <body bgcolor="#e7e8e9">
    <h1>Search Blocked</h1>
        <p><b>User:</b> <user/> </p>
       <p>Your search results have been blocked because your search settings are not in accordance with company
policy. In order to continue, please update your search settings so that Safe Search is set to the strictest
setting. If you are currently logged into your account, please also lock Safe Search and try your search
again.</p>
        <p>For more information, please refer to: <a href="<ssurl/>"><ssurl/></a></p>
        <p id="java_off"> Please enable JavaScript in your browser.<br></p>
        <p><b>Please contact your system administrator if you believe this message is in error.</b></p>


</body>
<script>

    // Grab the URL that's in the browser.
    var s_u = location.href;

    //bing
    // Matches the forward slashes in the beginning, anything, then ".bing." then anything followed by a non
greedy slash. Hopefully the first forward slash.
    var b_a = /^.*\/\/(.+\.bing\..+?)\//.exec(s_u);

    if (b_a) {

        s_u = s_u + "&adlt=strict";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You are being redirected to a safer search!';

    }

    //google
    // Matches the forward slashes in the beginning, anything, then ".google." then anything followed by a non
greedy slash. Hopefully the first forward slash.
    var g_a = /^.*\/\/(.+\.google\..+?)\//.exec(s_u);

    if (g_a) {

     s_u = s_u.replace(/&safe=off/ig,"");
        s_u = s_u + "&safe=active";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You are being redirected to a safer search!';

    }

    //yahoo
    // Matches the forward slashes in the beginning, anything, then ".yahoo."" then anything followed by a non
greedy slash. Hopefully the first forward slash.
    var y_a = /^.*\/\/(.+\.yahoo\..+?)\//.exec(s_u);
    if (y_a) {

        s_u = s_u.replace(/&vm=p/ig,"");
        s_u = s_u + "&vm=r";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You are being redirected to a safer search!';
    }
    document.getElementById("java_off").innerHTML = ' ';

    </script>
</html>
```

 © Palo Alto Networks, Inc.

| Enable Transparent Safe Search Enforcement (Continued) | | |
|---|---|---|
| Step 5 | Import the edited URL Filtering Safe Search Block page onto the firewall. | 1. To import the edited block page, select **Device > Response Pages > URL Filtering Safe Search Block Page**.<br>2. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.<br>3. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select **shared** to make it available to all virtual systems.<br>4. Click **OK** to import the file. |
| Step 6 | Enable SSL Forward Proxy decryption.<br><br>Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings. | 1. Add a custom URL category for the search sites:<br>a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.<br>b. Enter a **Name** for the category, such as SearchEngineDecryption.<br>c. **Add** the following to the Sites list:<br>`www.bing.*`<br>`www.google.*`<br>`search.yahoo.*`<br>d. Click **OK** to save the custom URL category object.<br>2. Follow the steps to Configure SSL Forward Proxy.<br>3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**. |
| Step 7 | Save the configuration. | Click **Commit**. |

# Set Up the PAN-DB Private Cloud

Use the following section to deploy one or more M-500 appliances as a PAN-DB private cloud within your network or data center and Configure the Firewalls to Access the PAN-DB Private Cloud.

| Set up the PAN-DB Private Cloud | | |
|---|---|---|
| Step 1 | Rack mount the M-500 appliance. | Refer to the M-500 Hardware Reference Guide for instructions. |
| Step 2 | Register the M-500 appliance. | For instruction on registering the M-500 appliance, see Register the Firewall. |
| Step 3 | Perform Initial Configuration of the M-500 Appliance.<br> <br>The M-500 appliance in PAN-DB mode uses two ports- MGT (Eth0) and Eth1; Eth2 is not used in PAN-DB mode. The management port is used for administrative access to the appliance and for obtaining the latest content updates from the PAN-DB public cloud. For communication between the appliance (PAN-DB server) and the firewalls on the network, you can use the MGT port or Eth1. | 1. Connect to the M-500 appliance in one of the following ways:<br>• Attach a serial cable from a computer to the Console port on the M-500 appliance and connect using a terminal emulation software (9600-8-N-1).<br>• Attach an RJ-45 Ethernet cable from a computer to the MGT port on the M-500 appliance. From a browser, go to https://192.168.1.1.Enabling access to this URL might require changing the IP address on the computer to an address in the 192.168.1.0 network (for example, 192.168.1.2).<br>2. When prompted, log in to the appliance. Log in using the default username and password (admin/admin). The appliance will begin to initialize.<br>3. Configure an network access settings including the IP address for the MGT interface:<br>**set deviceconfig system ip-address <server-IP> netmask <netmask> default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>**<br>where <server-IP> is the IP address you want to assign to the management interface of the server, <netmask> is the subnet mask, <gateway-IP> is the IP address of the network gateway, and <DNS-IP> is the IP address of the primary DNS server.<br>4. Configure an network access settings including the IP address for the Eth1 interface:<br>**set deviceconfig system eth1 ip-address <server-IP> netmask <netmask> default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>**<br>where <server-IP> is the IP address you want to assign to the data interface of the server, <netmask> is the subnet mask, <gateway-IP> is the IP address of the network gateway, and <DNS-IP> is the IP address of the DNS server.<br>5. Save your changes to the PAN-DB server.<br>**commit** |

 © Palo Alto Networks, Inc.

| Set up the PAN-DB Private Cloud | |
|---|---|
| Step 4 Switch to PAN-DB private cloud mode. | 1. To switch to PAN-DB mode, use the CLI command:<br>**request system system-mode pan-url-db**<br>You can switch from Panorama mode to PAN-DB mode and back; and from Panorama mode to Log Collector mode and back. Switching directly from PAN-DB mode to Log Collector mode or vice versa is not supported. When switching operational mode, a data reset is triggered. With the exception of management access settings, all existing configuration and logs will be deleted on restart.<br>2. Use the following command to verify that the mode is changed:<br>**show pan-url-cloud-status**<br>hostname: M-500<br>ip-address: 1.2.3.4<br>netmask: 255.255.255.0<br>default-gateway: 1.2.3.1<br>ipv6-address: unknown<br>ipv6-link-local-address: fe80:00/64<br>ipv6-default-gateway:<br>mac-address: 00:56:90:e7:f6:8e<br>time: Mon Apr 27 13:43:59 2015<br>uptime: 10 days, 1:51:28<br>family: m<br>model: M-500<br>serial: 0073010000xxx<br>sw-version: 7.0.0<br>app-version: 492-2638<br>app-release-date: 2015/03/19 20:05:33<br>av-version: 0<br>av-release-date: unknown<br>wf-private-version: 0<br>wf-private-release-date: unknown<br>logdb-version: 7.0.9<br>platform-family: m<br>pan-url-db: 20150417-220<br>**system-mode: Pan-URL-DB**<br>operational-mode: normal<br>3. Use the following command to check the version of the cloud database on the M-500 appliance:<br>**show pan-url-cloud-status**<br>Cloud status: Up<br>URL database version: 20150417-220 |

| Set up the PAN-DB Private Cloud | |
|---|---|
| Step 5 Install content and database updates.<br> The appliance only stores the currently running version of the content and one earlier version. | Pick one of the following methods of installing the content and database updates:<br>• If the PAN-DB server has direct Internet access use the following commands:<br>a. To check whether a new version is published use:<br>`request pan-url-db upgrade check`<br>b. To check the version that is currently installed on your server use:<br>`request pan-url-db upgrade info`<br>c. To download and install the latest version:<br>– `request pan-url-db upgrade download latest`<br>– `request pan-url-db upgrade install <version latest \| file>`<br>d. To schedule the M-500 appliance to automatically check for updates:<br>`set deviceconfig system update-schedule pan-url-db recurring weekly action download-and-install day-of-week <day of week> at <hr:min>`<br>• If the PAN-DB server is offline, access the Palo Alto Networks Support portal to download and save the content updates to an SCP server on your network. You can then import and install the updates using the following commands:<br>• `scp import pan-url-db remote-port <port-number> from username@host:path`<br>• `request pan-url-db upgrade install file <filename>` |

 © Palo Alto Networks, Inc.

| Set up the PAN-DB Private Cloud | |
|---|---|
| Step 6 Set up administrative access to the PAN-DB private cloud.<br> <br>The appliance has a default `admin` account. Any additional administrative users that you create can either be superusers (with full access) or superusers with read-only access.<br>PAN-DB private cloud does not support the use of RADIUS VSAs. If the VSAs used on the firewall or Panorama are used for enabling access to the PAN-DB private cloud, an authentication failure will occur. | • To set up a local administrative user on the PAN-DB server:<br>a. **configure**<br>b. **set mgt-config users <username> permissions role-based <superreader \| superuser> yes**<br>c. **set mgt-config users <username> password**<br>Enter password:**xxxxx**<br>Confirm password:**xxxxx**<br>d. **commit**<br>• To set up an administrative user with RADIUS authentication:<br>a. Create RADIUS server profile.<br>**set shared server-profile radius <server_profile_name> server <server_name> ip-address <ip_address> port <port_no> secret <shared_password>**<br>b. Create authentication-profile.<br>**set shared authentication-profile <auth_profile_name> user-domain <domain_name_for_authentication> allow-list <all> method radius server-profile <server_profile_name>**<br>c. Attach the authentication-profile to the user.<br>**set mgt-config users <username> authentication-profile <auth_profile_name>**<br>d. Commit the changes.<br>**commit**<br>• To view the list of users:.<br>**show mgt-config users**<br>`users { admin { phash fnRL/G5lXVMug; permissions { role-based { superuser yes; } } } admin_user_2 { permissions { role-based { superreader yes; } } authentication-profile RADIUS; } }` |
| Step 7 Configure the Firewalls to Access the PAN-DB Private Cloud. | |

## Configure the Firewalls to Access the PAN-DB Private Cloud

When using the PAN-DB public cloud, each firewall accesses the PAN-DB servers in the AWS cloud to download the list of eligible servers to which it can connect for URL lookups. With the PAN-DB private cloud, you must configure the firewalls with a (static) list of your PAN-DB private cloud servers that will be used for URL lookups. The list can contain up to 20 entries; IPv4 addresses, IPv6 addresses, and FQDNs are supported. Each entry on the list— IP address or FQDN—must be assigned to the management port and/or eth1 of the PAN-DB server.

**Configure the Firewalls to Access the PAN-DB Private Cloud.**

Step 1 Pick one of the following options based on the PAN-OS version on the firewall.

a. For firewalls running PAN-OS 7.0, Access the PAN-OS CLI or the web interface on the firewall.

- Use the following CLI command to configure access to the private cloud:

  **`set deviceconfig setting pan-url-db cloud-static-list <IP addresses> enable`**
- Or, in the web interface for each firewall,
  1. Select **Device > Setup >Content-ID,** edit the URL Filtering section.
  2. Enter the **PAN-DB Server** IP address(es) or FQDN(s). The list must be comma separated.

b. For firewalls running PAN-OS 5.0, 6.0, or 6.1, use the following CLI command to configure access to the private cloud:

**`debug device-server pan-url-db cloud-static-list-enable <IP addresses> enable`**

Step 2 **Commit** your changes.

Step 3 To verify that the change is effective, use the following CLI command on the firewall:

- **`show url-cloud status`**

```
Cloud status:            Up
URL database version:    20150417-220
```

To delete the entries for the private PAN-DB servers, and allow the firewalls to connect to the PAN-DB public cloud, use the command:

**`set deviceconfig setting pan-url-db cloud-static-list <IP addresses> disable`**

When you delete the list of private PAN-DB servers, a re-election process is triggered on the firewall. The firewall first checks for the list of PAN-DB private cloud servers and when it cannot find one, the firewall accesses the PAN-DB servers in the AWS cloud to download the list of eligible servers to which it can connect.

 © Palo Alto Networks, Inc.

# URL Filtering Use Case Examples

The following use cases show how to use App-ID to control a specific set of web-based applications and how to use URL categories as match criteria in a policy. When working with App-ID, it is important to understand that each App-ID signature may have dependencies that are required to fully control an application. For example, with Facebook applications, the App-ID facebook-base is required to access the Facebook website and to control other Facebook applications. For example, to configure the firewall to control Facebook email, you would have to allow the App-IDs facebook-base and facebook-mail. As another example, if you search Applipedia (the App-ID database) for LinkedIn, you will see that in order to control LinkedIn mail, you need to apply the same action to both App-IDs: linkedin-base and linkedin-mail. To determine application dependencies for App-ID signatures, visit Applipedia, search for the given application, and then click the application for details.

The User-ID feature is required to implement policies based on users and groups and a Decryption policy is required to identify and control websites that are encrypted using SSL/TLS.

This section includes two uses cases:

- Use Case: Control Web Access
- Use Case: Use URL Categories for Policy Matching

## Use Case: Control Web Access

When using URL filtering to control user website access, there may be instances where granular control is required for a given website. In this use case, a URL filtering profile is applied to the security policy that allows web access for your users and the *social-networking* URL category is set to block, but the allow list in the URL profile is configured to allow the social networking site Facebook. To further control Facebook, the company policy also states that only marketing has full access to Facebook and all other users within the company can only read Facebook posts and cannot use any other Facebook applications, such as email, posting, chat, and file sharing. To accomplish this requirement, App-ID must be used to provide granular control over Facebook.

The first security rule will allow marketing to access the Facebook website as well as all Facebook applications. Because this allow rule will also allow access to the Internet, threat prevention profiles are applied to the rule, so traffic that matches the policy will be scanned for threats. This is important because the allow rule is terminal and will not continue to check other rules if there is a traffic match.

| **Control Web Access** | | |
|---|---|---|
| Step 1 | Confirm that URL filtering is licensed. | 1. Select **Device > Licenses** and confirm that a valid date appears for the URL filtering database that will used. This will either be PAN-DB or BrightCloud.<br>2. If a valid license is not installed, see Enable PAN-DB URL Filtering. |
| Step 2 | Confirm that User-ID is working. User-ID is required to create policies based on users and groups. | 1. To check Group Mapping from the CLI, enter the following command:<br>`show user group-mapping statistics`<br>2. To check User Mapping from the CLI, enter the following command:<br>`show user ip-user-mapping-mp all`<br>3. If statistics do not appear and/or IP address to user mapping information is not displayed, see User-ID. |
| Step 3 | Set up a URL filtering profile by cloning the default profile. | 1. Select **Objects > Security Profiles > URL Filtering** and select the **default** profile.<br>2. Click the **Clone** icon. A new profile should appear named **default-1**.<br>3. Select the new profile and rename it. |

 © Palo Alto Networks, Inc.

| Control Web Access (Continued) | | |
|---|---|---|
| Step 4 | Configure the URL filtering profile to block social-networking and allow Facebook. | 1. Modify the new URL filtering profile and in the **Category** list scroll to **social-networking** and in the **Action** column click on **allow** and change the action to **block**.<br>2. In the **Allow List**, enter `facebook.com`, press enter to start a new line and then type `*.facebook.com`. Both of these formats are required, so all URL variants a user may use will be identified, such as facebook.com, www.facebook.com, and https://facebook.com.<br> 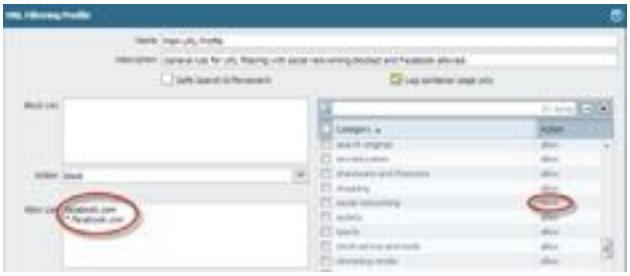 <br>3. Click **OK** to save the profile. |
| Step 5 | Apply the new URL filtering profile to the security policy rule that allows web access from the user network to the Internet. | 1. Select **Policies > Security** and click on the policy rule that allows web access.<br>2. On the **Actions** tab, select the URL profile you just created from the **URL Filtering** drop-down.<br> 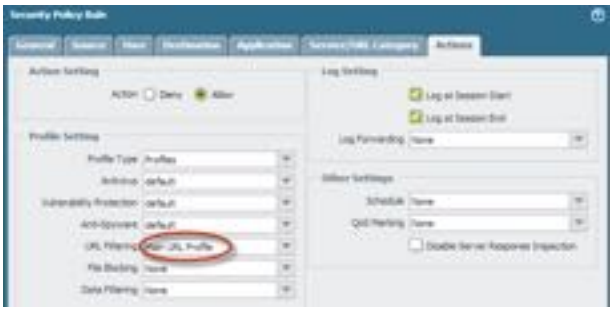 <br>3. Click **OK** to save. |

| Control Web Access (Continued) | |
|---|---|
| Step 6 Create the security policy rule that will allow marketing access the Facebook website and all Facebook applications.<br><br>This rule must precede other rules because:<br>• It is a specific rule. More specific rules must precede other rules.<br>• Allow rule will terminate when a traffic match occurs. | 1. Select **Policies > Security** and click **Add**.<br>2. Enter a **Name** and optionally a **Description** and **Tag**(s).<br>3. On the **Source** tab add the zone where the users are connected.<br>4. On the **User** tab in the **Source User** section click **Add**.<br>5. Select the directory group that contains your *marketing* users.<br>6. On the **Destination** tab, select the zone that is connected to the Internet.<br>7. On the **Applications** tab, click **Add** and add the *facebook* App-ID signature.<br>8. On the **Actions** tab, add the default profiles for **Antivirus**, **Vulnerability Protection**, and **Anti-Spyware**.<br> <br>9. Click **OK** to save the security profile.<br>The *facebook* App-ID signature used in this policy rule encompasses all Facebook applications, such as facebook-base, facebook-chat, and facebook-mail, so this is the only App-ID signature required in this rule.<br>With this rule in place, when a marketing employee attempts to access the Facebook website or any Facebook application, the rule matches based on the user being part of the marketing group. For traffic from any user outside of marketing, the rule will be skipped because there would not be a traffic match and rule processing would continue. |


© Palo Alto Networks, Inc.

**Control Web Access (Continued)**

Step 7 Configure the security policy to block all other users from using any Facebook applications other than simple web browsing. The easiest way to do this is to clone the marketing allow policy and then modify it.

1. From **Policies > Security** click the marketing Facebook allow policy you created earlier to highlight it and then click the **Clone** icon.
2. Enter a **Name** and optionally enter a **Description** and **Tag**('s).
3. On the **User** tab highlight the marketing group and delete it and in the drop-down select **any**.
4. On the **Applications** tab, click the *facebook* App-ID signature and delete it.
5. Click **Add** and add the following App-ID signatures:
   - facebook-apps
   - facebook-chat
   - facebook-file-sharing
   - facebook-mail
   - facebook-posting
   - facebook-social-plugin
6. On the **Actions** tab in the **Action Setting** section, select **Deny**. The profile settings should already be correct because this rule was cloned.

7. Click **OK** to save the security profile.
8. Ensure that this new deny rule is listed after the marketing allow rule, to ensure that rule processing occurs in the correct order to allow marketing users and then to deny/limit all other users.
9. Click **Commit** to save the configuration.

With these security policy rules in place, any user who is part of the marketing group will have full access to all Facebook applications and any user that is not part of the marketing group will only have read-only access to the Facebook website and will not be able to use Facebook applications such as post, chat, email, and file sharing.

## Use Case: Use URL Categories for Policy Matching

URL categories can also be used as match criteria in the following policy types: Captive Portal, Decryption, Security, and QoS. In this use case, URL categories will be used in Decryption policy rules to control which web categories should be decrypted or not decrypted. The first rule is a no-decrypt rule that will not decrypt user traffic if the website category is *financial-services* or *health-and-medicine* and the second rule will decrypt all other traffic. The decryption policy type is *ssl-forward-proxy*, which is used for controlling decryption for all outbound connections performed by users.

**Configure a Decryption Policy Based on URL Category**

Step 1 Create the no-decrypt rule that will be listed first in the decryption policies list. This will prevent any website that is in the *financial-services* or *health-and-medicine* URL categories from being decrypted.

1. Select **Policies > Decryption** and click **Add**.
2. Enter a **Name** and optionally enter a **Description** and **Tag**(s).
3. On the **Source** tab, add the zone where the users are connected.
4. On the **Destination** tab, enter the zone that is connected to the Internet.
5. On the **URL Category** tab, click **Add** and select the *financial-services* and *health-and-medicine* URL categories.
6. On the **Options** tab, set the action to **No Decrypt**.
7. (Optional) Although the firewall does not decrypt and inspect the traffic for the session, you can attach a **Decryption profile** if you want to enforce the server certificates used during the session. The decryption profile allows you to configure the firewall to terminate the SSL connection either when the server certificates are expired or when the server certificates are issues by an untrusted issuer.

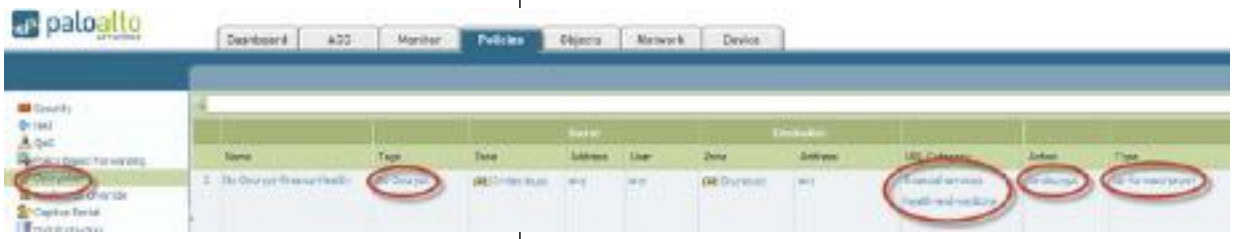


8. Click **OK** to save the policy rule.

 © Palo Alto Networks, Inc.

| Configure a Decryption Policy Based on URL Category (Continued) | | |
|---|---|---|
| Step 2 | Create the decryption policy rule that will decrypt all other traffic. | 1. Select the no-decrypt policy you created previously and then click **Clone**.<br>2. Enter a **Name** and optionally enter a **Description** and **Tag**(s).<br>3. On the **URL Category** tab, select *financial-services* and *health-and-medicine* and then click the **Delete** icon.<br>4. On the **Options** tab, set the action to **Decrypt** and the **Type** to **SSL Forward Proxy**.<br>5. (Optional) Attach a **Decryption profile** to specify the server certificate verification, unsupported mode checks and failure checks for the SSL traffic. See Configure SSL Forward Proxy for more details.<br>6. Ensure that this new decryption rule is listed after the no-decrypt rule to ensure that rule processing occurs in the correct order, so websites in the *financial-services* and *health-and-medicine* are not decrypted<br>7. Click **OK** to save the policy rule. |
| Step 3 | (BrightCloud only) Enable cloud lookups for dynamically categorizing a URL when the category is not available on the local database on the firewall. | 1. Access the CLI on the firewall.<br>2. Enter the following commands to enable Dynamic URL Filtering:<br>a. **configure**<br>b. **set deviceconfig setting url dynamic-url yes**<br>c. **commit** |
| Step 4 | Save the configuration. | Click **Commit**. |

With these two decrypt policies in place, any traffic destined for the *financial-services* or *health-and-medicine* URL categories will not be decrypted. All other traffic will be decrypted.

Now that you have a basic understanding of the powerful features of URL filtering, App-ID, and User-ID, you can apply similar policies to your firewall to control any application in the Palo Alto Networks App-ID signature database and control any website contained in the URL filtering database.

For help in troubleshooting URL filtering issues, see Troubleshoot URL Filtering.

# Troubleshoot URL Filtering

The following topics provide troubleshooting guidelines for diagnosing and resolving common URL filtering problems.


- Problems Activating PAN-DB
- PAN-DB Cloud Connectivity Issues
- URLs Classified as Not-Resolved
- Incorrect Categorization
- URL Database Out of Date


## Problems Activating PAN-DB

The following table describes procedures that you can use to resolve issues with activating PAN-DB.

| Troubleshoot PAN-DB Activation Issues |
|---|
| 1. Access the PAN-OS CLI. |
| 2. Verify whether PAN-DB has been activated by running the following command:<br>`admin@PA-200> show system setting url-database`<br>If the response is `paloaltonetworks`, then PAN-DB is the active vendor. |
| 3. Verify that the firewall has a valid PAN-DB license by running the following command:<br>`admin@PA-200> request license info`<br>You should see the license entry Feature: PAN_DB URL Filtering. If the license is not installed, you will need to obtain and install a license. See Configure URL Filtering. |
| 4. After the license is installed, download a new PAN-DB seed database by running the following command:<br>`admin@PA-200> request url-filtering download paloaltonetworks region` *<region>* |
| 5. Check the download status by running the following command:<br>`admin@PA-200> request url-filtering download status vendor paloaltonetworks`<br>• If the message is different from `PAN-DB download: Finished successfully`, stop here; there may be a problem connecting to the cloud. Attempt to solve the connectivity issue by performing basic network troubleshooting between the firewall and the Internet. For more information, see PAN-DB Cloud Connectivity Issues.<br>• If the message is `PAN-DB download: Finished successfully`, the firewall successfully downloaded the URL seed database. Try to enable PAN-DB again by running the following command:<br>`admin@PA-200> set system setting url-database paloaltonetworks` |
| 6. If the problems persists, contact Palo Alto Networks support. |


© Palo Alto Networks, Inc.

## PAN-DB Cloud Connectivity Issues

To check cloud connectivity, run the following command:

```
admin@pa-200> show url-cloud status
```

If the cloud is accessible, the expected response is similar to the following:

```
admin@PA-200> show url-cloud status
PAN-DB URL Filtering
License :                          valid
Current cloud server :             s0000.urlcloud.paloaltonetworks.com
Cloud connection :                 connected
URL database version - device :    2013.11.18.000
URL database version - cloud :     2013.11.18.000  ( last update time
2013/11/19
 13:20:51 )
URL database status :              good
URL protocol version - device :    pan/0.0.2
URL protocol version - cloud :     pan/0.0.2
Protocol compatibility status :    compatible
```

If the cloud is note accessible, the expected response is similar to the following:

```
admin@PA-200> show url-cloud status
PAN-DB URL Filtering
License :                          valid
Cloud connection :                 not connected
URL database version - device :    2013.11.18.000
URL database version - cloud :     2013.11.18.000  ( last update time
2013/11/19
 13:20:51 )
URL database status :              good
URL protocol version - device :    pan/0.0.2
URL protocol version - cloud :     pan/0.0.2
Protocol compatibility status :    compatible
```

The following table describes procedures that you can use to resolve issues based on the output of the `show url-cloud status` command, how to ping the URL cloud servers, and what to check if the firewall is in a High Availability (HA) configuration.

| Troubleshoot Cloud Connectivity Issues |
| --- |
| • PAN-DB URL Filtering license field shows invalid—Obtain and install a valid PAN-DB license.<br>• URL database status is out of date—Download a new seed database by running the following command:<br>`admin@pa-200> request url-filtering download paloaltonetworks region <region>`<br>• URL protocol version shows not compatible—Upgrade PAN-OS to the latest version.<br>• Attempt to ping the PAN-DB cloud server from the firewall by running the following command:<br>`admin@pa-200> ping source <ip-address> host s0000.urlcloud.paloaltonetworks.com`<br>For example, if your management interface IP address is 10.1.1.5, run the following command:<br>`admin@pa-200> ping source 10.1.1.5 host s0000.urlcloud.paloaltonetworks.com`<br>• If the firewall is in an HA configuration, verify that the HA state of the devices supports connectivity to the cloud systems. You can determine the HA state by running the following command:<br>`admin@pa-200> show high-availability state`<br>Connection to the cloud will be blocked if the firewall is not in one of the following states:<br>  • active<br>  • active-primary<br>  • active-secondary<br>If the problem persists, contact Palo Alto Networks support. |

## URLs Classified as Not-Resolved

The following table describes procedures you can use to resolve issues where some or all of the URLs being identified by PAN-DB are classified as *Not-resolved*:

| Troubleshoot URLs Classified as Not-Resolved |
| --- |
| 1. Check the PAN-DB cloud connection by running the following command:<br>`admin@PA-200> show url-cloud status`<br>The Cloud connection: field should show *connected*. If you see anything other than *connected*, any URL that do not exist in the management plane cache will be categorized as *not-resolved*. To resolve this issue, see PAN-DB Cloud Connectivity Issues. |
| 2. If the cloud connection status shows *connected*, check the current utilization of the firewall. If the firewall's performance is spiking, URL requests may be dropped (may not reach the management plane), and will be categorized as *not-resolved*.<br>To view system resources, run the following command and view the *%CPU* and *%MEM* columns:<br>`admin@PA-200> show system resources`<br>You can also view system resources from the firewall's web interfaces by clicking the **Dashboard** tab and viewing the **System Resources** section. |
| 3. If the problem persist, contact Palo Alto Networks support. |

 © Palo Alto Networks, Inc.

## Incorrect Categorization

The following steps describe the procedures you can use if you identify a URL that does not have the correct categorization. For example, if the URL *paloaltonetworks.com* was categorized as *alcohol-and-tobacco*, the categorization is not correct; the category should be *computer-and-internet-info*.

**Troubleshoot Incorrect Categorization Issues**

1. Verify the category in the dataplane by running the following command:

   `admin@PA-200>` show running url <*URL*>

   For example, to view the category for the Palo Alto Networks website, run the following command:

   ```
   admin@PA-200> show running url paloaltonetworks.com
   ```

   If the URL stored in the dataplane cache has the correct category (computer-and-internet-info in this example), then the categorization is correct and no further action is required. If the category is not correct, continue to the next step.

2. Verify if the category in the management plane by running the command:

   `admin@PA-200> test url-info-host` <*URL*>

   For example:

   ```
   admin@PA-200> test url-info-host paloaltonetworks.com
   ```

   If the URL stored in the management plane cache has the correct category, remove the URL from the dataplane cache by running the following command:

   `admin@PA-200> clear url-cache url` <*URL*>

   The next time the device requests the category for this URL, the request will be forwarded to the management plane. This will resolve the issue and no further action is required. If this does not solve the issue, go to the next step to check the URL category on the cloud systems.

3. Verify the category in the cloud by running the following command:

   `admin@PA-200> test url-info-cloud <URL>`

4. If the URL stored in the cloud has the correct category, remove the URL from the dataplane and the management plane caches.

   Run the following command to delete a URL from the dataplane cache:

   ```
   admin@PA-200> clear url-cache url <URL>
   ```

   Run the following command to delete a URL from the management plane cache:

   ```
   admin@PA-200> delete url-database url <URL>
   ```

   The next time the device queries for the category of the given URL, the request will be forwarded to the management plane and then to the cloud. This should resolve the category lookup issue. If problems persist, see the next step to submit a categorization change request.

5. To submit a change request from the web interface, go to the URL log and select the log entry for the URL you would like to have changed.

**Troubleshoot Incorrect Categorization Issues**

6. Click the **Request Categorization** change link and follow instructions. You can also request a category change from the Palo Alto Networks Test A Site website by searching for the URL and then clicking the **Request Change** icon. To view a list of all available categories with descriptions of each category, refer to https://urlfiltering.paloaltonetworks.com/CategoryList.aspx.

   If your change request is approved, you will receive an email notification. You then have two options to ensure that the URL category is updated on the firewall:

   - Wait until the URL in the cache expires and the next time the URL is accessed by a user, the new categorization update will be put in the cache.
   - Run the following command to force an update in the cache:

     ```
     admin@PA-200> request url-filtering update url <URL>
     ```

## URL Database Out of Date

If you have observed through the syslog or the CLI that PAN-DB is out-of-date, it means that the connection from the firewall to the URL Cloud is blocked. This usually occurs when the URL database on the firewall is too old (version difference is more than three months) and the cloud cannot update the firewall automatically. In order to resolve this issue, you will need to re-download an initial seed database from the cloud (this operation is not blocked). This will result in an automatic re-activation of PAN-DB.

To manually update the database, perform one of the following steps:

- From the web interface, select **Device > Licenses** and in the **PAN-DB URL Filtering** section click the **Re-Download** link.
- From the CLI, run the following command:

  ```
  admin@PA-200> request url-filtering download paloaltonetworks region <region_name>
  ```

When the seed database is re-download, the URL cache in the management plane and dataplane will be purged. The management plane cache will then be re-populated with the contents of the new seed database.

 © Palo Alto Networks, Inc.

# Quality of Service

Quality of Service (QoS) is a set of technologies that work on a network to guarantee its ability to dependably run high-priority applications and traffic under limited network capacity. QoS technologies accomplish this by providing differentiated handling and capacity allocation to specific flows in network traffic. This enables the network administrator to assign the order in which traffic is handled, and the amount of bandwidth afforded to traffic.

Palo Alto Networks Application Quality of Service (QoS) provides basic QoS applied to networks and extends it to provide QoS to applications and users.

Use the following topics to learn about and configure Palo Alto Networks Application QoS:

- QoS Overview
- QoS Concepts
- Configure QoS
- Configure QoS for a Virtual System
- Enforce QoS Based on DSCP Classification
- QoS Use Cases

- Use the Palo Alto Networks product comparison tool to view the QoS features supported on your firewall platform. Select two or more product platforms and click **Compare Now** to view QoS feature support for each platform (for example, you can check if your firewall platform supports QoS on subinterfaces and if so, the maximum number of subinterfaces on which QoS can be enabled).
- QoS on Aggregate Ethernet (AE) interfaces is supported on PA-7000 Series, PA-5000 Series, PA-3000 Series, and PA-2000 Series firewalls running PAN-OS 7.0 or later release versions.

# QoS Overview

Use QoS to prioritize and adjust quality aspects of network traffic. You can assign the order in which packets are handled and allot bandwidth, ensuring preferred treatment and optimal levels of performance are afforded to selected traffic, applications, and users.

Service quality measurements subject to a QoS implementation are bandwidth (maximum rate of transfer), throughput (actual rate of transfer), latency (delay), and jitter (variance in latency). The capability to shape and control these service quality measurements makes QoS of particular importance to high-bandwidth, real-time traffic such as voice over IP (VoIP), video conferencing, and video-on-demand that has a high sensitivity to latency and jitter. Additionally, use QoS to achieve outcomes such as the following:

- Prioritize network and application traffic, guaranteeing high priority to important traffic or limiting non-essential traffic.
- Achieve equal bandwidth sharing among different subnets, classes, or users in a network.
- Allocate bandwidth externally or internally or both, applying QoS to both upload and download traffic or to only upload or download traffic.
- Ensure low latency for customer and revenue-generating traffic in an enterprise environment.
- Perform traffic profiling of applications to ensure bandwidth usage.

QoS implementation on a Palo Alto Networks firewall begins with three primary configuration components that support a full QoS solution: a QoS Profile, a QoS Policy, and setting up the QoS Egress Interface. Each of these options in the QoS configuration task facilitate a broader process that optimizes and prioritizes the traffic flow and allocates and ensures bandwidth according to configurable parameters.

The figure QoS Traffic Flow shows traffic as it flows from the source, is shaped by the firewall with QoS enabled, and is ultimately prioritized and delivered to its destination.

**QoS Traffic Flow**

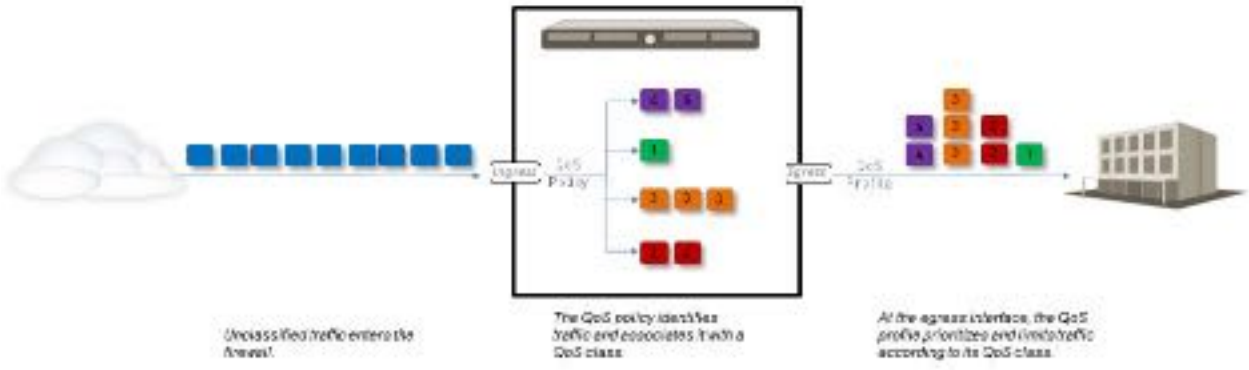

The QoS configuration options allow you to control the traffic flow and define it at different points in the flow. The QoS Traffic Flow indicates where the configurable options define the traffic flow. Use the QoS Profile to define QoS classes and use the QoS Policy to associate QoS classes with selected traffic. Enable the QoS Profile on an interface to shape traffic according to the QoS configuration as it flows through the network.

 © Palo Alto Networks, Inc.

You can configure a QoS Profile and QoS Policy individually or in any order, according to your preference. Each of the QoS configuration options has components that influence the definition of the other options and the QoS configuration options can be used to create a full and granular QoS policy or can be used sparingly with minimal administrator action.

Each firewall model supports a maximum number of ports that can be configured with QoS. Refer to the spec sheet for your firewall model or use the product comparison tool to view QoS feature support for two or more firewalls on a single page.

# QoS Concepts

Use the following topics to learn about the different components and mechanisms of a QoS configuration on a Palo Alto Networks firewall:

- QoS for Applications and Users
- QoS Profile
- QoS Classes
- QoS Policy
- QoS Egress Interface
- QoS Clear Text and Tunneled Traffic

## QoS for Applications and Users

A Palo Alto Networks firewall provides basic QoS, controlling traffic leaving the firewall according to network or subnet, and extends the power of QoS to also classify and shape traffic according to application and user. The Palo Alto Networks firewall provides this capability by integrating the features App-ID and User-ID with the QoS configuration. App-ID and User-ID entries that exist to identify specific applications and users in your network are available in the QoS configuration so that you can easily specify applications and users to apply QoS to.

You can use a QoS Policy in the web interface (**Policies** > **QoS**) to apply QoS specifically to an application's traffic:

Or to a user's traffic:

See App-ID and User-ID for more information on these features.


© Palo Alto Networks, Inc.

## QoS Profile

Use a QoS profile to define values of up to eight QoS classes contained within that single profile (**Network** > **Network Profiles** > **QoS Profile**):

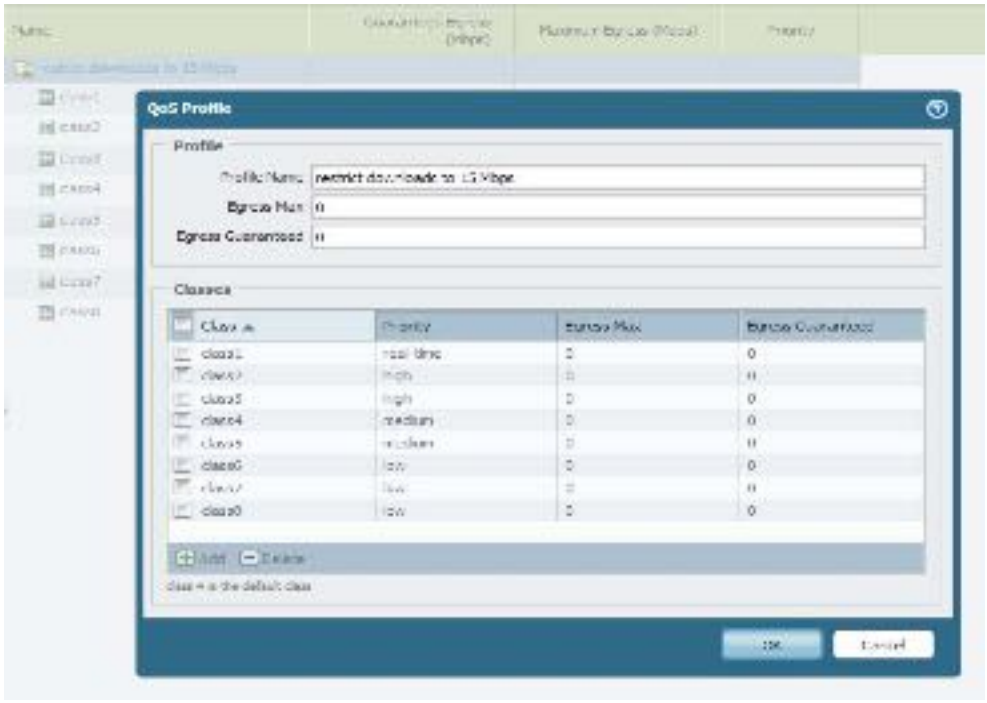


You enable QoS by applying a QoS profile to the egress interface for network, application, and user traffic, and for clear text or tunneled traffic. An interface configured with QoS shapes traffic according to the QoS profile class definitions and the traffic associated with those classes in the QoS policy.

A default QoS Profile is available on the firewall. The default profile and the classes defined in the profile do not have predefined maximum or guaranteed bandwidth limits.

You can set bandwidth limits for a QoS profile and/or set limits for individual QoS classes within the QoS profile. The total guaranteed bandwidth limits of all eight QoS classes in a QoS Profile cannot exceed the total bandwidth allocated to that QoS Profile. Enabling QoS on a physical interface includes setting the maximum bandwidth for traffic leaving the firewall through this interface. A QoS profile's guaranteed bandwidth (the **Egress Guaranteed** field) should not exceed the bandwidth allocated to the physical interface that QoS is enabled on.

For details, see Add a QoS profile.

## QoS Classes

A QoS class determines the priority and bandwidth for traffic it is assigned to. In the web interface, use the QoS profile to define QoS classes (**Network** > **Network Profiles** > **QoS Profile**):

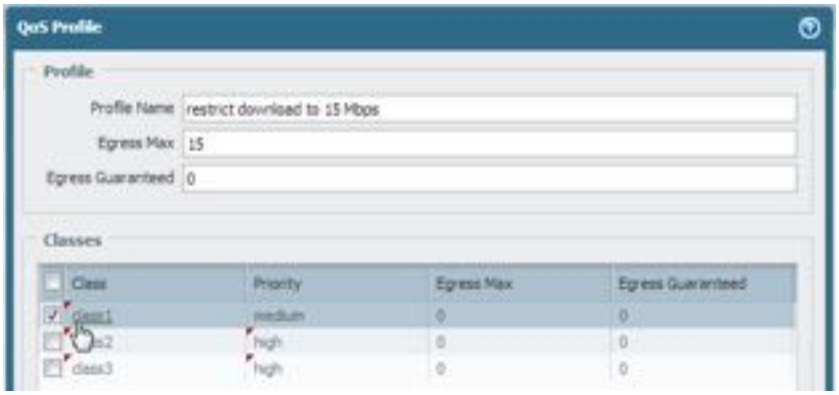


Defining a QoS class includes setting the class's Priority, maximum bandwidth (Egress Max), and guaranteed bandwidth (Egress Guaranteed).

Real-time priority is typically used for applications that are particularly sensitive to latency, such as voice and video applications.

Use the QoS policy to assign a QoS class to specified traffic (**Policies** > **QoS**):

There are up to eight definable QoS classes in a single QoS profile. Unless otherwise configured, traffic that does not match a QoS class is assigned a class of 4.

QoS priority queuing and bandwidth management, the fundamental mechanisms of a QoS configuration, are configured within the QoS class definition (see Step 3). Queuing priority is determined by the priority set for a QoS class. Bandwidth management is determined according to the maximum and guaranteed bandwidths set for a QoS class.

 © Palo Alto Networks, Inc.

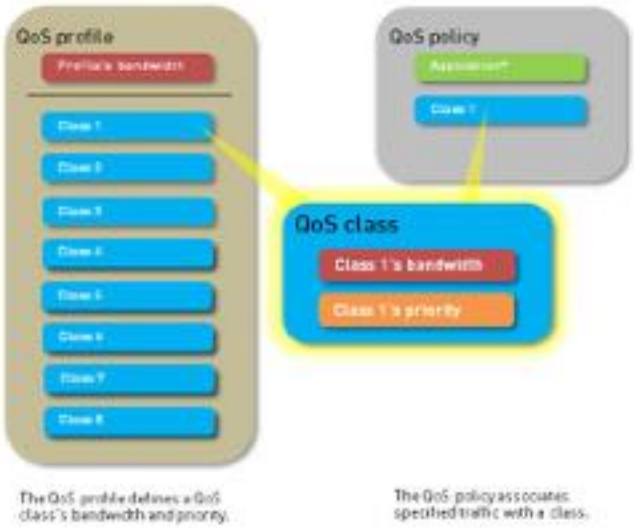


The queuing and bandwidth management mechanisms determine the order of traffic and how traffic is handled upon entering or leaving a network:

- **QoS Priority:** One of four QoS priorities can be defined in a QoS class: real-time, high, medium, and low. When a QoS class is associated with specific traffic, the priority defined in that QoS class is assigned to the traffic. Packets in the traffic flow are then queued according to their priority until the network is ready to process them. This method of priority queuing provides the capability to ensure that important traffic, applications, or users takes precedence.
- **QoS Class Bandwidth Management:** QoS class bandwidth management provides the capability to control traffic flows on a network so that traffic does not exceed network capacity, resulting in network congestion, or to allocate specific bandwidth limits to traffic, applications, or users. You can set overall limits on bandwidth using the QoS profile or set limits for individual QoS classes. A QoS profile and QoS classes in the profile have guaranteed and maximum bandwidth limits. The guaranteed bandwidth limit (Egress Guaranteed) ensures that any amount of traffic up to that set bandwidth limit is processed. The maximum bandwidth limit (Egress Max) sets the total limit of bandwidth allocated to either the QoS Profile or QoS Class. Traffic in excess of the Maximum Bandwidth limit is dropped. The total bandwidth limits and guaranteed bandwidth limits of QoS classes in a QoS profile cannot exceed the bandwidth limit of the QoS profile.

## QoS Policy

Use a QoS policy rule to define traffic to receive QoS treatment (either preferential treatment or bandwidth-limiting) and assigns such traffic a QoS class of service.

Define a QoS policy rule to match to traffic based on:

- Applications and application groups.
- Source zones, source addresses, and source users.
- Destination zones and destination addresses.

- Services and service groups limited to specific TCP and/or UDP port numbers.
- URL categories, including custom URL categories.
- Differentiated Services Code Point (DSCP) and Type of Service (ToS) values, which are used to indicate the level of service requested for traffic, such as high priority or best effort delivery.

Set up multiple QoS policy rules (**Policies** > **QoS**) to associate different types of traffic with different QoS Classes of service.

## QoS Egress Interface

Enabling a QoS profile on the egress interface of the traffic identified for QoS treatment completes a QoS configuration. The ingress interface for QoS traffic is the interface on which the traffic enters the firewall. The egress interface for QoS traffic is the interface that traffic leaves the firewall from. QoS is always enabled and enforced on a traffic flow's egress interface. The egress interface in a QoS configuration can either be the external- or internal-facing interface of the firewall, depending on the flow of the traffic receiving QoS treatment.

For example, in an enterprise network, if you are limiting employees' download traffic from a specific website, the egress interface in the QoS configuration is the firewall's internal interface, as the traffic flow is from the Internet, through the firewall, and to your company network. Alternatively, when limiting employees' upload traffic to the same website, the egress interface in the QoS configuration is the firewall's external interface, as the traffic you are limiting flows from your company network, through the firewall, and then to the Internet.

See Step 3 to learn how to Identify the egress interface for applications that you identified as needing QoS treatment.

 © Palo Alto Networks, Inc.

## QoS Clear Text and Tunneled Traffic

At the minimum, enabling a QoS interfaces requires you to select a default QoS profile to be used to shape all clear text traffic egressing the interface. However, when setting up or modifying a QoS interface, granular QoS settings can be applied to outgoing clear text traffic and tunneled traffic. QoS preferential treatment and bandwidth limiting can be enforced for individual tunnel interfaces and/or for clear text traffic originating from different source interfaces and source subnets.

On Palo Alto Networks firewalls, *tunneled traffic* refers to tunnel interface traffic, specifically IPSec traffic in tunnel mode.

# Configure QoS

Follow these steps to configure Quality of Service (QoS), which includes creating a QoS profile, creating a QoS policy, and enabling QoS on an interface.

| Configure QoS | | |
|---|---|---|
| Step 1 | Identify the traffic to which to apply QoS.<br><br>This example shows how to use QoS to limit web browsing. | Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.<br><br>Click any application name to display detailed application information. |
| Step 2 | Identify the egress interface for applications that you identified as needing QoS treatment.<br><br>The egress interface for traffic depends on the traffic flow. If you are shaping incoming traffic, the egress interface is the internal-facing interface. If you are shaping outgoing traffic, the egress interface is the external-facing interface. | Select **Monitor > Logs > Traffic** to view the device's traffic logs.<br><br>To filter and only show logs for a specific application:<br>• If an entry is displayed for the application, click the underlined link in the Application column then click the Submit icon.<br>• If an entry is not displayed for the application, click the Add Log icon and search for the application.<br><br>The **Egress I/F** in the traffic logs displays each application's egress interface. To display the **Egress I/F** column if it is not displayed by default:<br>• Click any column header to add a column to the log:<br> 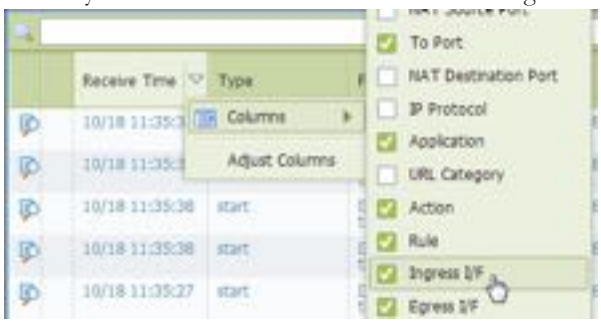 <br>• Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface listed in the Destination section:<br> <br>In this example, the egress interface for web-browsing traffic is ethernet 1/1. |

 © Palo Alto Networks, Inc.

**Configure QoS (Continued)**

Step 3 Add a QoS profile.

A QoS profile allows you to define the eight classes of service that traffic can receive, including the priority and bandwidth for each class, as well as the total, combined bandwidth for all eight classes.

You can edit any existing QoS profile, including the default, by clicking the QoS profile name.

1. Select **Network > Network Profiles > QoS Profile** and click **Add** to open the **QoS Profile** dialog.
2. Enter a descriptive **Profile Name**.
3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS Profile.
4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS Profile.

   Any traffic that exceeds the **Egress Guaranteed** value is best effort and not guaranteed.
5. In the Classes section, specify how to treat up to eight individual QoS classes:
   a. Click **Add** to add a class to the QoS Profile.
   b. Select the **Priority** for the class.
   c. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.
   d. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.
6. Click **OK** to save the QoS Profile.

In the following example, the QoS Profile named Limit Web Browsing limits traffic identified as Class 2 traffic to maximum bandwidth of 50Mbps and a guaranteed bandwidth of 2Mbps. Any traffic that is associated with class 2 in a the QoS policy (Step 4) is subject to these limits.

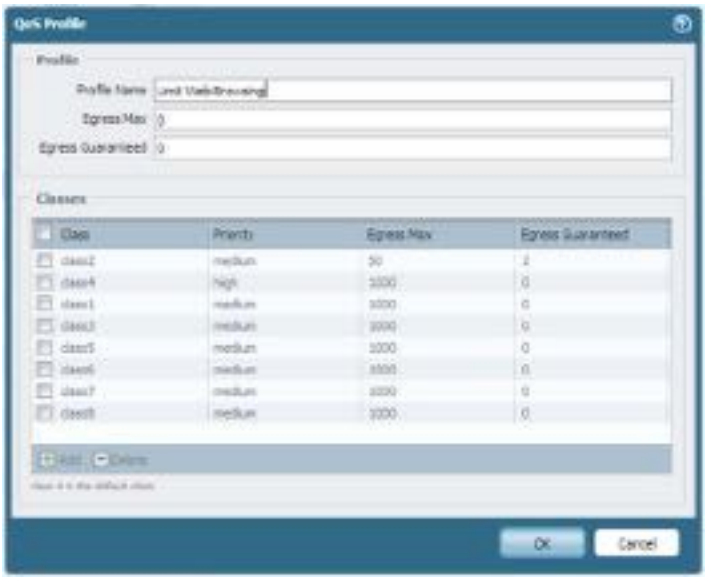

| Configure QoS (Continued) | |
|---|---|
| Step 4 Add a QoS policy rule.<br><br>A QoS policy allows define traffic to receive QoS treatment, and to assign traffic matched to the policy rule a QoS class of service. | 1. Select **Policies > QoS** and click **Add** to open the QoS Policy Rule dialog.<br>2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**.<br>3. Specify the traffic to which the QoS Policy Rule will apply. Use the **Source**, **Destination**, **Application**, **Service/URL Category**, and **DSCP/ToS** tabs to define matching parameters for identifying traffic (use the **DSCP/ToS** tab to Enforce QoS Based on DSCP Classification).<br>For example, select the **Application** tab, click **Add**, and select web-browsing to apply the QoS Policy Rule to that application:<br> 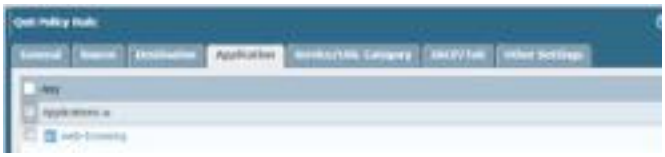 <br>(Optional) Define additional parameters. For example, on the **Source** tab, click **Add** to limit web-browsing for a specific user, in this case, user1:<br> <br>4. On the **Other Settings** tab, select a QoS Class to assign to the QoS Policy Rule. For example, assign Class 2 to the user1's web-browsing traffic:<br> <br>5. Click **OK** to save the QoS Policy Rule. |


© Palo Alto Networks, Inc.

**Configure QoS (Continued)**

Step 5 Enable the QoS Profile on a physical interface.

You can configure settings to select clear text and tunneled traffic for unique QoS treatment, in addition to the QoS configuration on the physical interface:

- To define QoS treatment for clear text traffic based on source interface and subnet, perform step 4.
- To define QoS treatment for a specific tunnel interface(s), perform step 5.

See QoS Clear Text and Tunneled Traffic for details on enforcing QoS for clear text and tunneled traffic.

Check if the platform you're using supports enabling QoS on a subinterface by reviewing a summary of the Product Specifications.

It is a best practice to always define the **Egress Max** value for a QoS interface.

1. Select **Network > QoS** and click **Add** to open the QoS Interface dialog.
2. Enable QoS on the physical interface:
   a. On the **Physical Interface** tab, select the **Interface Name** of the interface to which to enable QoS.

      In the example, Ethernet 1/1 is the egress interface for web-browsing traffic (see Step 2).

   b. Select **Turn on QoS feature on this interface**.
3. On the **Physical Interface** tab, select a QoS profile to apply by default to all **Clear Text** traffic.

   (Optional) Use the Tunnel Interface field to apply a QoS profile by default to all tunneled traffic.

For example, enable QoS on ethernet 1/1 and apply the QoS Profile named Limit Web Browsing as the default QoS Profile for clear text traffic.

4. (Optional) Select **Clear Text Traffic** and configure more granular QoS settings for clear text traffic:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
   - Click **Add** to further specify QoS treatment for clear text traffic, based on source interface and source subnet.
5. (Optional) Select **Tunneled Traffic** and configure more granular QoS settings for tunnel interfaces:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
   - Click **Add** and attach a QoS profile to a single tunnel interface.
6. Click **OK** and **Commit** the changes to enable QoS on the interface.

| Configure QoS (Continued) | |
|---|---|
| Step 6 Verify QoS configuration. | Select **Network > QoS** and then **Statistics** to view QoS bandwidth, active sessions of a selected QoS class, and active applications for the selected QoS class.<br>For example, see the statistics for ethernet 1/1 with QoS enabled:<br> 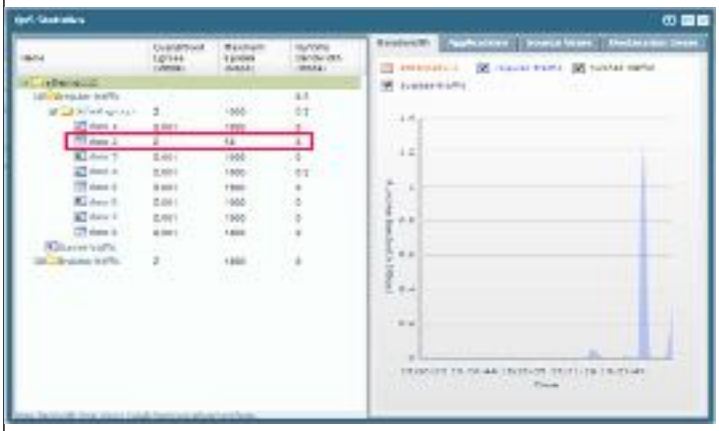 <br>Class 2 traffic limited to 2Mbps of guaranteed bandwidth and a maximum bandwidth of 50Mbps.<br>Continue to click the tabs to display further information regarding applications, source users, destination users, security rules and QoS rules.<br>Bandwidth limits shown on the **QoS Statistics** window include a hardware adjustment factor. |


© Palo Alto Networks, Inc.

# Configure QoS for a Virtual System

QoS can be configured for a single or several virtual systems configured on a Palo Alto Networks firewall. Because a virtual system is an independent firewall, QoS must be configured independently for a single virtual system.

Configuring QoS for a virtual system is similar to configuring QoS on a physical firewall, with the exception that configuring QoS for a virtual system requires specifying the source and destination of traffic. Because a virtual system exists without set physical boundaries and because traffic in a virtual environment spans more than one virtual system, specifying source and destination zones and interfaces for traffic is necessary to control and shape traffic for a single virtual system.

The example below shows two virtual systems configured on firewall. VSYS 1 (purple) and VSYS 2 (red) each have QoS configured to prioritize or limit two distinct traffic flows, indicated by their corresponding purple (VSYS 1) and red (VSYS 2) lines. The QoS nodes indicate the points at traffic is matched to a QoS policy and assigned a QoS class of service, and then later indicate the point at which traffic is shaped as it egresses the firewall.

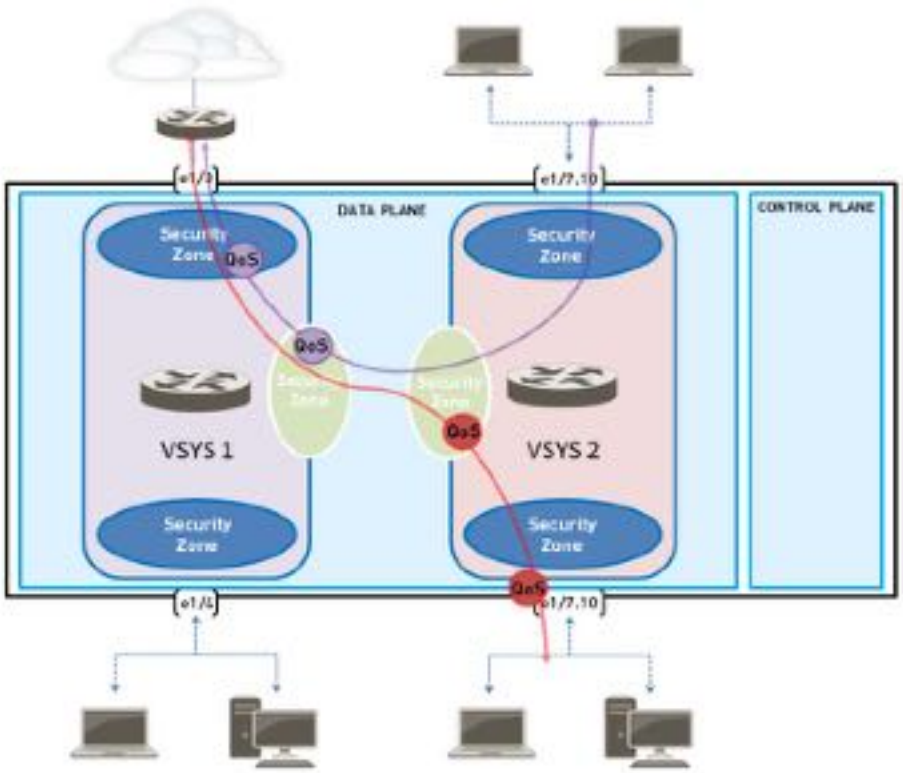


Refer to the Virtual Systems (VSYS) tech note for information on Virtual Systems and how to configure them.

| Configure QoS in a Virtual System Environment | | |
|---|---|---|
| Step 1 | Confirm that the appropriate interfaces, virtual routers, and security zones are associated with each virtual system. | • To view configured interfaces, select **Network > Interface**.<br>• To view configured zones, select **Network > Zones**.<br>• To view information on defined virtual routers, select **Network > Virtual Routers**. |

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 2 Identify traffic to apply QoS to. | Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.<br>To view information for a specific virtual system, select the virtual system from the **Virtual System** drop-down:<br> 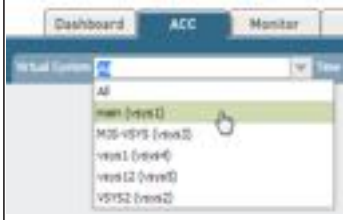 <br><br>Click any application name to display detailed application information. |

 © Palo Alto Networks, Inc.

**Configure QoS in a Virtual System Environment**

Step 3 Identify the egress interface for applications that you identified as needing QoS treatment.

In a virtual system environment, QoS is applied to traffic on the traffic's egress point on the virtual system. Depending the configuration and QoS policy for a virtual system, the egress point of QoS traffic could be associated with a physical interface or could be a zone.

This example shows how to limit web-browsing traffic on vsys 1.

Select **Monitor > Logs > Traffic** to view traffic logs. Each entry has the option to display columns with information necessary to configure QoS in a virtual system environment:

- virtual system
- egress interface
- ingress interface
- source zone
- destination zone

To display a column if it is not displayed by default:

- Click any column header to add a column to the log:

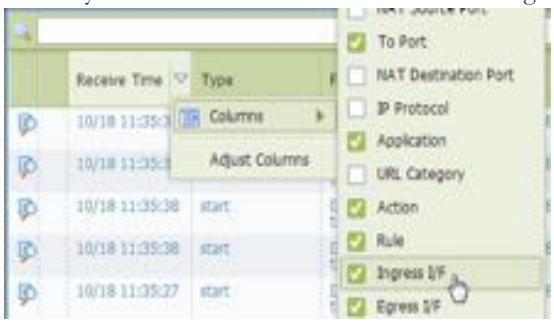


- Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface, as well as source and destination zones, in the **Source** and **Destination** sections:

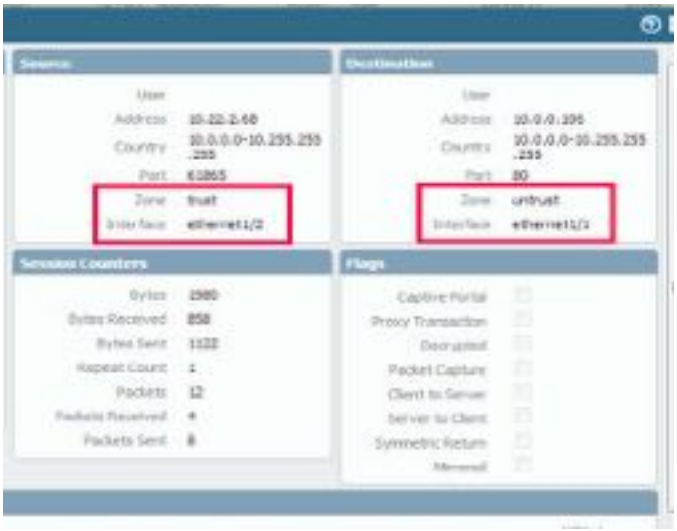

For example, for web-browsing traffic from VSYS 1, the ingress interface is ethernet 1/2, the egress interface is ethernet 1/1, the source zone is *trust* and the destination zone is *untrust*.

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 4 Create a QoS Profile.<br><br>You can edit any existing QoS Profile, including the default, by clicking the profile name. | 1. Select **Network > Network Profiles > QoS Profile** and click **Add** to open the QoS Profile dialog.<br>2. Enter a descriptive **Profile Name**.<br>3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS profile.<br>4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS profile.<br><br> Any traffic that exceeds the QoS profile's egress guaranteed limit is best effort but is not guaranteed.<br>5. In the Classes section of the **QoS Profile**, specify how to treat up to eight individual QoS classes:<br>a. Click **Add** to add a class to the QoS Profile.<br>b. Select the **Priority** for the class.<br>c. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.<br>d. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.<br>6. Click **OK** to save the QoS profile. |

 © Palo Alto Networks, Inc.

**Configure QoS in a Virtual System Environment**

Step 5 Create a QoS policy.

In an environment with multiple virtual systems, traffic spans more than one virtual system. Because of this, when you are enabling QoS for a virtual system, you must define traffic to receive QoS treatment based on source and destination zones. This ensures that the traffic is prioritized and shaped only for that virtual system (and not for other virtual systems through which the traffic might flow).

1. Select **Policies > QoS** and **Add** a QoS Policy Rule.
2. Select **General** and give the QoS Policy Rule a descriptive **Name**.
3. Specify the traffic to which the QoS policy rule will apply. Use the **Source**, **Destination**, **Application**, and **Service/URL Category** tabs to define matching parameters for identifying traffic.

   For example, select **Application** and **Add** web-browsing to apply the QoS policy rule to that application:

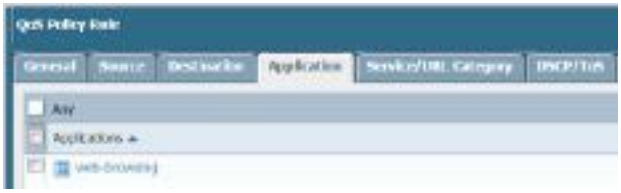


4. Select **Source** and **Add** the source zone of vsys 1 web-browsing traffic.

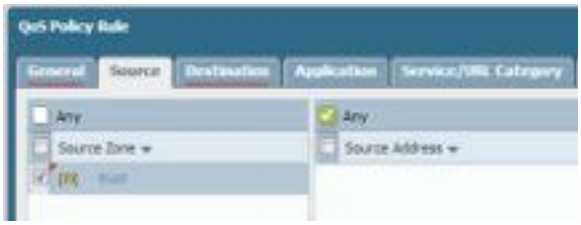


5. Select **Destination** and **Add** the destination zone of vsys 1 web-browsing traffic.

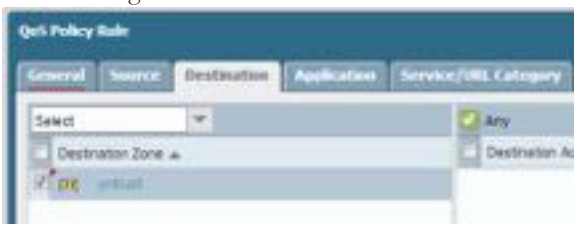


6. Select **Other Settings** and select a **QoS Class** to assign to the QoS policy rule. For example, assign Class 2 to web-browsing traffic on vsys 1:

7. Click **OK** to save the QoS policy rule.

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 6 Enable the QoS Profile on a physical interface.<br><br>It is a best practice to always define the **Egress Max** value for a QoS interface. | 1. Select **Network > QoS** and click **Add** to open the QoS Interface dialog.<br>2. Enable QoS on the physical interface:<br>a. On the **Physical Interface** tab, select the **Interface Name** of the interface to apply the QoS Profile to.<br>In this example, ethernet 1/1 is the egress interface for web-browsing traffic on vsys 1 (see Step 2).<br> <br>b. Select **Turn on QoS feature on this interface**.<br>3. On the **Physical Interface** tab, select the default QoS profile to apply to all **Clear Text** traffic.<br>(Optional) Use the **Tunnel Interface** field to apply a default QoS profile to all tunneled traffic.<br>4. (Optional) On the **Clear Text Traffic** tab, configure additional QoS settings for clear text traffic:<br>• Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.<br>• Click **Add** to apply a QoS Profile to selected clear text traffic, further selecting the traffic for QoS treatment according to source interface and source subnet (creating a QoS node).<br>5. (Optional) On the **Tunneled Traffic** tab, configure additional QoS settings for tunnel interfaces:<br>• Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.<br>• Click **Add** to associate a selected tunnel interface with a QoS Profile.<br>6. Click **OK** to save changes.<br>7. **Commit** the changes. |


© Palo Alto Networks, Inc.

| Configure QoS in a Virtual System Environment | |
|---|---|
| Step 7 Verify QoS configuration. | • Select **Network > QoS** to view the QoS Policies page. The **QoS Policies** page verifies that QoS is enabled and includes a **Statistics** link. Click the Statistics link to view QoS bandwidth, active sessions of a selected QoS node or class, and active applications for the selected QoS node or class.<br>• In a multi-vsys environment, sessions cannot span multiple systems. Multiple sessions are created for one traffic flow if the traffic passes through more than one virtual system. To browse sessions running on the firewall and view applied QoS Rules and QoS Classes, select **Monitor > Session Browser**. |

# Enforce QoS Based on DSCP Classification

A Differentiated Services Code Point (DSCP) is a packet header value that can be used to indicate the level of service requested for traffic, such as high priority or best effort delivery. Session-Based DSCP Classification allows you to both honor the service class requested for traffic and to mark a session to receive continued QoS treatment. Session-based DSCP extends the power of Quality of Service (QoS), which polices traffic as it passes through the firewall, by allowing all network devices between the firewall and the client to also police traffic. All inbound and outbound traffic for a session can receive continuous QoS/DSCP treatment as it flows through your network. For example, inbound return traffic from an external server can now be treated with the same QoS/DSCP priority that the firewall initially enforced for the outbound flow. Network devices between the firewall and end user will also then enforce the same priority for the return traffic (and any other outbound or inbound traffic for the session).

Use the following steps to enable Session-Based DSCP Classification. Start by configuring QoS based on DSCP marking detected at the beginning of a session. You can then continue to enable the firewall to mark the return flow for a session with the same DSCP value used to enforce QoS for the initial outbound flow.

| **Provide QoS Based on DSCP/ToS Marking** | |
|---|---|
| Before you set up or modify a QoS policy to control traffic based on DSCP/ToS code point values, make sure that you have performed the preliminary steps for configuring QoS. | The following steps are useful before setting up a QoS policy:<br>❑ Identify the traffic to which to apply QoS.<br>❑ Identify the egress interface for traffic to be enforced with QoS.<br>See Configure QoS for details on performing these steps. |

 © Palo Alto Networks, Inc.

**Provide QoS Based on DSCP/ToS Marking (Continued)**

Step 1 Apply QoS to traffic based on the DSCP value detected at the beginning of the session.

1. Select **Policies > QoS** and **Add** or modify an existing QoS rule and populate required fields.
2. **Add** a **DSCP/ToS** rule and give the rule a descriptive **Name**. You can choose to add multiple DSCP/ToS rules to a single QoS rule to enforce the same QoS priority for sessions with different DSCP values.
3. Select the **Type** of DSCP/ToS marking for the QoS rule to match to traffic:

   It is a best practice to use a single DSCP type to manage and prioritize your network traffic.

   - **Expedited Forwarding (EF)**: Can be used to request low loss, low latency and guaranteed bandwidth for traffic. Packets with EF codepoints are typically guaranteed highest priority delivery.
   - **Assured Forwarding (AF)**: Can be used to provide reliable delivery for applications. Packets with AF codepoint indicate a request for the traffic to receive higher priority treatment than best effort service provides (though packets with an EF codepoint will continue to take precedence over those with an AF codepoint).
   - **Class Selector (CS)**: Can be used to provide backward compatibility with network devices that use the IP precedence field to mark priority traffic.
   - **IP Precedence (ToS)**: Can be used by legacy network devices to mark priority traffic (the IP Precedence header field was used to indicate the priority for a packet before the introduction of the DSCP classification).
   - **Custom Codepoint**: Create a custom codepoint to match to traffic by entering a **Codepoint Name** and **Binary Value**.

   For example, select the **Assured Forwarding (AF)** to ensure traffic marked with an **AF** codepoint value has higher priority for reliable delivery over applications marked to receive lower priority.
4. Match the QoS policy to traffic on a more granular scale by specifying the **Codepoint** value. For example, with Assured Forwarding (AF) selected as the **Type** of DSCP value for the policy to match, further specify an AF **Codepoint** value such as AF11.

   When Expedited Forwarding (EF) is selected as the **Type** of DSCP marking, a granular **Codepoint** value cannot be specified. The QoS policy will match to traffic marked with any EF codepoint value.
5. Select **Other Settings** and a **QoS Class** to assign to traffic matched to the QoS rule. For example, assign Class 1 to sessions where a DSCP marking of AF11 is detected for the first packet in the session.
6. Click **OK** to save the QoS rule.

| | | |
|---|---|---|
| Step 2 | Define the QoS priority for traffic to receive when it is matched to a QoS rule based the DSCP marking detected at the beginning of a session. | 1. Select **Network > Network Profiles > QoS Profile** and **Add** or modify an existing QoS profile. For details on profile options to set priority and bandwidth for traffic, see QoS Concepts and Configure QoS.<br>2. **Add** or modify a profile class. For example, because Step 1 showed steps to classify AF11 traffic as Class 1 traffic, you could add or modify a **class1** entry.<br>3. Select a **Priority** for the class of traffic, such as **high**.<br>4. Click **OK** to save the QoS Profile. |

| Provide QoS Based on DSCP/ToS Marking (Continued) | | |
|---|---|---|
| Step 3 | Enable QoS on an interface. | Select **Network > QoS** and **Add** or modify an existing interface and **Turn on QoS feature on this interface**.<br>In this example, traffic with an AF11 DSCP marking is matched to the QoS rule and assigned Class 1. The QoS profile enabled on the interface enforces high priority treatment for Class 1 traffic as it egresses the firewall (the session *outbound* traffic). |
| Step 4 | Enable DSCP Marking.<br>Mark return traffic with a DSCP value, enabling the inbound flow for a session to be marked with the same DSCP value detected for the outbound flow. | 1. Select **Policies > Security** and **Add** or modify a security policy.<br>2. Select **Actions** and in the **QoS Marking** drop-down, choose **Follow-Client-to-Server-Flow**.<br>3. Click **OK** to save your changes.<br>Completing this step enables the firewall to mark traffic with the same DSCP value that was detected at the beginning of a session (in this example, the firewall would mark return traffic with the DSCP AF11 value). While configuring QoS allows you to shape traffic as it egresses the firewall, enabling this option in a security rule allows the other network devices intermediate to the firewall and the client to continue to enforce priority for DSCP marked traffic. |
| Step 5 | Save the configuration. | **Commit** your changes. |

 © Palo Alto Networks, Inc.

# QoS Use Cases

The following use cases demonstrate how to use QoS in common scenarios:

- Use Case: QoS for a Single User
- Use Case: QoS for Voice and Video Applications

© Palo Alto Networks, Inc.

## Use Case: QoS for a Single User

A CEO finds that during periods of high network usage, she is unable to access enterprise applications to respond effectively to critical business communications. The IT admin wants to ensure that all traffic to and from the CEO receives preferential treatment over other employee traffic so that she is guaranteed not only access to, but high performance of, critical network resources.

### Apply QoS to a Single User

Step 1 The admin creates the QoS profile *CEO_traffic* to define how traffic originating from the CEO will be treated and shaped as it flows out of the company network:

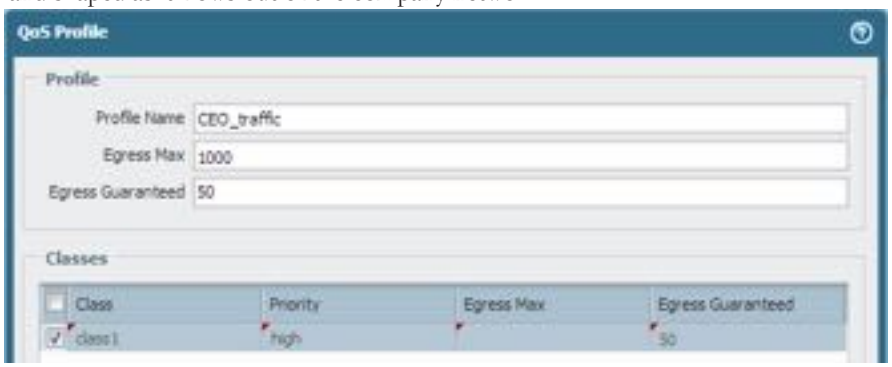


The admin assigns a guaranteed bandwidth (**Egress Guaranteed**) of 50 Mbps to ensure that the CEO will have that amount that bandwidth guaranteed to her at all times (more than she would need to use), regardless of network congestion.

The admin continues by designating Class 1 traffic as high priority and sets the profile's maximum bandwidth usage (**Egress Max**) to 1000 Mbps, the same maximum bandwidth for the interface that the admin will enable QoS on. The admin is choosing to not restrict the CEO's bandwidth usage in any way.

It is a best practice to populate the **Egress Max** field for a QoS profile, even if the max bandwidth of the profile matches the max bandwidth of the interface. The QoS profile's max bandwidth should never exceed the max bandwidth of the interface you are planning to enable QoS on.

 © Palo Alto Networks, Inc.

**Apply QoS to a Single User (Continued)**

Step 2 The admin creates a QoS policy to identify the CEO's traffic (**Policies** > **QoS**) and assigns it the class that he defined in the QoS profile (see Step 1). Because User-ID is configured, the admin uses the **Source** tab in the QoS policy to singularly identify the CEO's traffic by her company network username. (If User-ID is not configured, the administrator could **Add** the CEO's IP address under **Source Address**. See User-ID.):

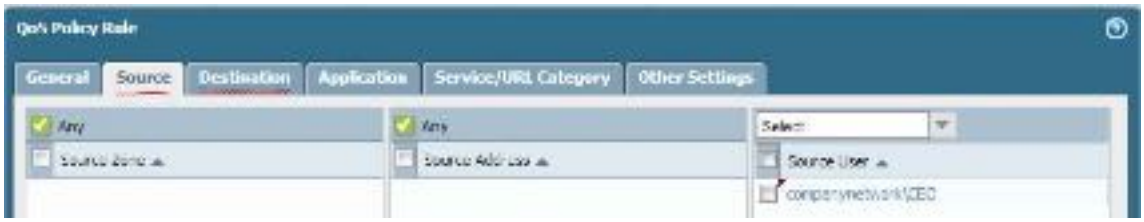


The admin associates the CEO's traffic with Class 1 (**Other Settings** tab) and then continues to populate the remaining required policy fields; the admin gives the policy a descriptive **Name** (**General** tab) and selects **Any** for the **Source Zone** (**Source** tab) and **Destination Zone** (**Destination** tab):

Step 3 Now that Class 1 is associated with the CEO's traffic, the admin enables QoS by checking **Turn on QoS feature on interface** and selecting the traffic flow's egress interface. The egress interface for the CEO's traffic flow is the external-facing interface, in this case, ethernet 1/2:

Because the admin wants to ensure that all traffic originating from the CEO is guaranteed by the QoS profile and associated QoS policy he created, he selects the *CEO_traffic* to apply to **Clear Text** traffic flowing from ethernet 1/2.

**Apply QoS to a Single User (Continued)**

Step 4 After committing the QoS configuration, the admin navigates to the **Network > QoS** page to confirm that the QoS profile *CEO_traffic* is enabled on the external-facing interface, ethernet 1/2:

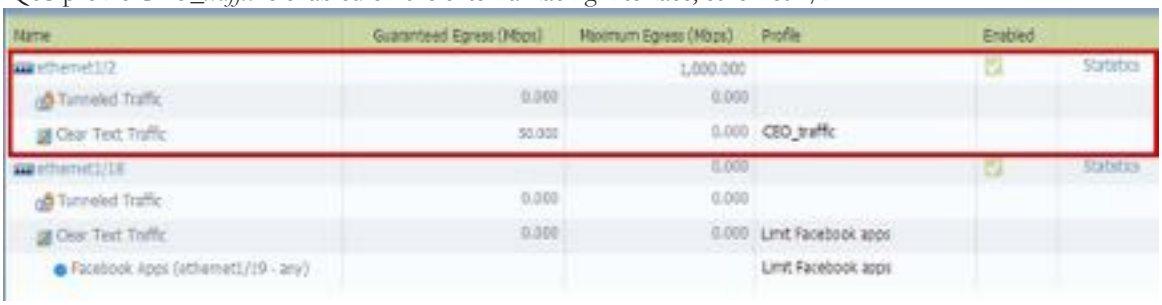


He clicks **Statistics** to view how traffic originating with the CEO (Class 1) is being shaped as it flows from ethernet 1/2:

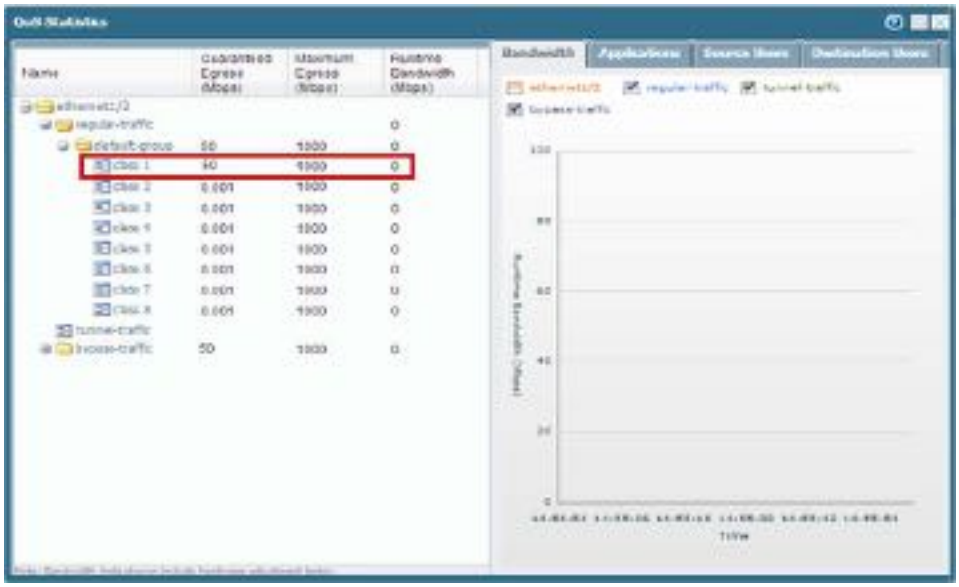


This case demonstrates how to apply QoS to traffic originating from a single source user. However, if you also wanted to guarantee or shape traffic to a destination user, you could configure a similar QoS setup. Instead of, or in addition to this work flow, create a QoS policy that specifies the user's IP address as the **Destination Address** on the **Policies > QoS** page (instead of specifying the user's source information, as shown in Step 2) and then enable QoS on the network's internal-facing interface on the **Network > QoS** page (instead of the external-facing interface, as shown in Step 3.)


© Palo Alto Networks, Inc.

## Use Case: QoS for Voice and Video Applications

Voice and video traffic is particularly sensitive to measurements that the QoS feature shapes and controls, especially latency and jitter. For voice and video transmissions to be audible and clear, voice and video packets cannot be dropped, delayed, or delivered inconsistently. A best practice for voice and video applications, in addition to guaranteeing bandwidth, is to guarantee priority to voice and video traffic.

In this example, employees at a company branch office are experiencing difficulties and unreliability in using video conferencing and Voice over IP (VoIP) technologies to conduct business communications with other branch offices, with partners, and with customers. An IT admin intends to implement QoS in order to address these issues and ensure effective and reliable business communication for the branch employees. Because the admin wants to guarantee QoS to both incoming and outgoing network traffic, he will enable QoS on both the firewall's internal- and external-facing interfaces.

**Ensure Quality for Voice and Video Applications**

Step 1 The admin creates a QoS profile, defining Class 2 so that any traffic associated with Class 2 receives real-time priority and on an interface with a maximum bandwidth of 1000 Mbps, is guaranteed a bandwidth of 250 Mbps at all times, including peak periods of network usage.

Real-time priority is typically recommended for applications affected by latency, and is particularly useful in guaranteeing performance and quality of voice and video applications.

On the **Network > Network Profiles > Qos Profile** page, the admin clicks **Add**, enters the **Profile Name** *ensure voip-video traffic* and defines Class 2 traffic.

**Ensure Quality for Voice and Video Applications (Continued)**

Step 2 The admin creates a QoS policy to identify voice and video traffic. Because the company does not have one standard voice and video application, the admin wants to ensure QoS is applied to a few applications that are widely and regularly used by employees to communicate with other offices, with partners, and with customers. On the **Policies > QoS > QoS Policy Rule > Applications** tab, the admin clicks **Add** and opens the **Application Filter** window. The admin continues by selecting criteria to filter the applications he wants to apply QoS to, choosing the Subcategory *voip-video*, and narrowing that down by specifying only voip-video applications that are both low-risk and widely-used.

The application filter is a dynamic tool that, when used to filter applications in the QoS policy, allows QoS to be applied to all applications that meet the criteria of *voip-video*, *low risk*, and *widely used* at any given time.

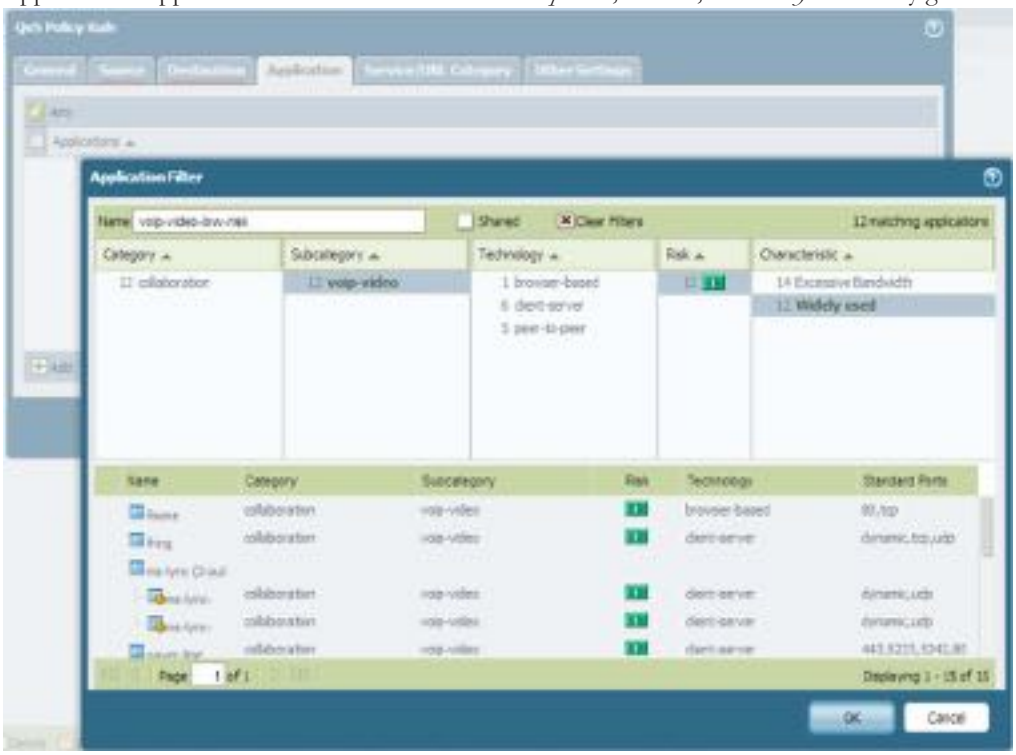


The admin names the **Application Filter** *voip-video-low-risk* and includes it in the QoS policy:

The admin names the QoS policy *Voice-Video* and associates the *voip-video-low-risk* application filter with Class 2 traffic (as he defined it in Step 1). He is going to use the *Voice-Video* QoS policy for both incoming and outgoing QoS traffic, so he sets **Source** and **Destination** information to **Any**:

 © Palo Alto Networks, Inc.

**Ensure Quality for Voice and Video Applications (Continued)**

Step 3 Because the admin wants to ensure QoS for both incoming and outgoing voice and video communications, he enables QoS on the network's external-facing interface (to apply QoS to outgoing communications) and to the internal-facing interface (to apply QoS to incoming communications).

The admin begins by enabling the QoS profile he created in Step 1, *ensure voice-video traffic* (Class 1 in this profile is associated with policy created in Step 2, *Voice-Video*) on the external-facing interface, in this case, ethernet 1/2.

He then enables the same QoS profile *ensure voip-video traffic* on the internal-facing interface, in this case, ethernet 1/1.

Step 4 The admin confirms that QoS is enabled for both incoming and outgoing voice and video traffic:

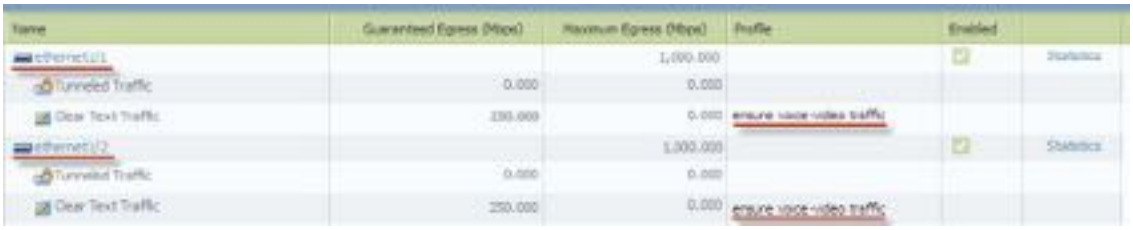


The admin has successfully enabled QoS on both the network's internal- and external-facing interfaces. Real-time priority is now ensured for voice and video application traffic as it flows both into and out of the network, ensuring that these communications, which are particularly sensitive to latency and jitter, can be used reliably and effectively to perform both internal and external business communications.

 © Palo Alto Networks, Inc.

# VPNs

Virtual private networks (VPNs) create tunnels that allow users/systems to connect securely over a public network, as if they were connecting over a local area network (LAN). To set up a VPN tunnel, you need a pair of devices that can authenticate each other and encrypt the flow of information between them. The devices can be a pair of Palo Alto Networks firewalls, or a Palo Alto Networks firewall along with a VPN-capable device from another vendor.


- VPN Deployments
- Site-to-Site VPN Overview
- Site-to-Site VPN Concepts
- Set Up Site-to-Site VPN
- Site-to-Site VPN Quick Configs


© Palo Alto Networks, Inc.

# VPN Deployments

The Palo Alto Networks firewall supports the following VPN deployments:

- **Site-to-Site VPN**— A simple VPN that connects a central site and a remote site, or a hub and spoke VPN that connects a central site with multiple remote sites. The firewall uses the IP Security (IPSec) set of protocols to set up a secure tunnel for the traffic between the two sites. See Site-to-Site VPN Overview.
- **Remote User-to-Site VPN**—A solution that uses the GlobalProtect agent to allow a remote user to establish a secure connection through the firewall. This solution uses SSL and IPSec to establish a secure connection between the user and the site. Refer to the GlobalProtect Administrator's Guide.
- **Large Scale VPN**— The Palo Alto Networks GlobalProtect Large Scale VPN (LSVPN) provides a simplified mechanism to roll out a scalable hub and spoke VPN with up to 1024 satellite offices. The solution requires Palo Alto Networks firewalls to be deployed at the hub and at every spoke. It uses certificates for device authentication, SSL for securing communication between all components, and IPSec to secure data. See Large Scale VPN (LSVPN).

**Figure: VPN Deployments**

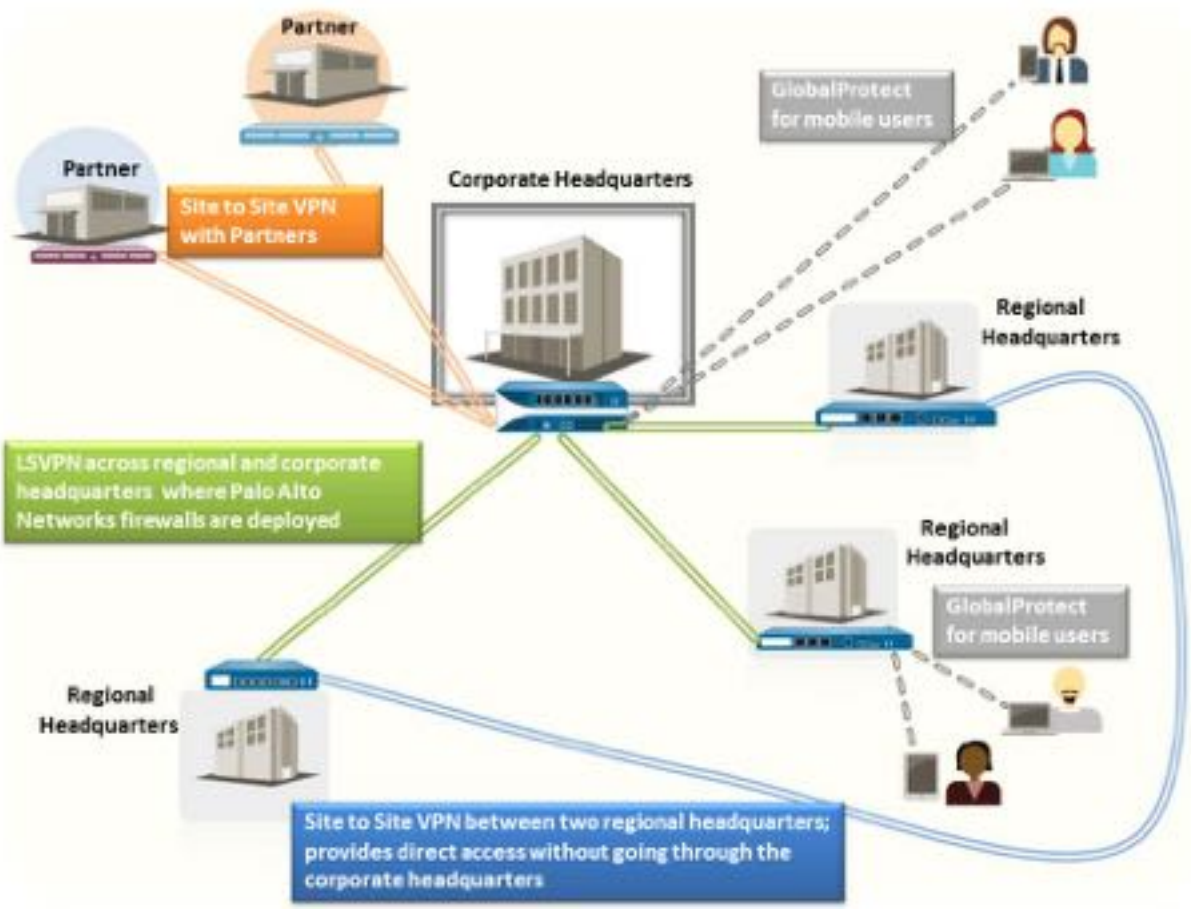


 © Palo Alto Networks, Inc.

# Site-to-Site VPN Overview

A VPN connection that allows you to connect two Local Area Networks (LANs) is called a site-to-site VPN. You can configure route-based VPNs to connect Palo Alto Networks firewalls located at two sites or to connect a Palo Alto Networks firewall with a third-party security device at another location. The firewall can also interoperate with third-party policy-based VPN devices; the Palo Alto Networks firewall supports route-based VPN.

The Palo Alto Networks firewall sets up a route-based VPN, where the firewall makes a routing decision based on the destination IP address. If traffic is routed to a specific destination through a VPN tunnel, then it is handled as VPN traffic.

The IP Security (IPSec) set of protocols is used to set up a secure tunnel for the VPN traffic, and the information in the TCP/IP packet is secured (and encrypted if the tunnel type is ESP). The IP packet (header and payload) is embedded in another IP payload, and a new header is applied and then sent through the IPSec tunnel. The source IP address in the new header is that of the local VPN peer and the destination IP address is that of the VPN peer on the far end of the tunnel. When the packet reaches the remote VPN peer (the firewall at the far end of the tunnel), the outer header is removed and the original packet is sent to its destination.

In order to set up the VPN tunnel, first the peers need to be authenticated. After successful authentication, the peers negotiate the encryption mechanism and algorithms to secure the communication. The Internet Key Exchange (IKE) process is used to authenticate the VPN peers, and IPSec Security Associations (SAs) are defined at each end of the tunnel to secure the VPN communication. IKE uses digital certificates or preshared keys, and the Diffie Hellman keys to set up the SAs for the IPSec tunnel.The SAs specify all of the parameters that are required for secure transmission— including the security parameter index (SPI), security protocol, cryptographic keys, and the destination IP address—encryption, data authentication, data integrity, and endpoint authentication.

The following figure shows a VPN tunnel between two sites. When a client that is secured by VPN Peer A needs content from a server located at the other site, VPN Peer A initiates a connection request to VPN Peer B. If the security policy permits the connection, VPN Peer A uses the IKE Crypto profile parameters (IKE phase 1) to establish a secure connection and authenticate VPN Peer B. Then, VPN Peer A establishes the VPN tunnel using the IPSec Crypto profile, which defines the IKE phase 2 parameters to allow the secure transfer of data between the two sites.

**Figure: Site-to-Site VPN**

# Site-to-Site VPN Concepts

A VPN connection provides secure access to information between two or more sites. In order to provide secure access to resources and reliable connectivity, a VPN connection needs the following components:

- IKE Gateway
- Tunnel Interface
- Tunnel Monitoring
- Internet Key Exchange (IKE) for VPN
- IKEv2

## IKE Gateway

The Palo Alto Networks firewalls or a firewall and another security device that initiate and terminate VPN connections across the two networks are called the IKE Gateways. To set up the VPN tunnel and send traffic between the IKE Gateways, each peer must have an IP address—static or dynamic—or FQDN. The VPN peers use preshared keys or certificates to mutually authenticate each other.

The peers must also negotiate the mode—main or aggressive—for setting up the VPN tunnel and the SA lifetime in IKE Phase 1. Main mode protects the identity of the peers and is more secure because more packets are exchanged when setting up the tunnel. Main mode is the recommended mode for IKE negotiation if both peers support it. Aggressive mode uses fewer packets to set up the VPN tunnel and is hence faster but a less secure option for setting up the VPN tunnel.

See Set Up an IKE Gateway for configuration details.

## Tunnel Interface

To set up a VPN tunnel, the Layer 3 interface at each end must have a logical *tunnel* interface for the firewall to connect to and establish a VPN tunnel. A tunnel interface is a logical (virtual) interface that is used to deliver traffic between two endpoints. Each tunnel interface can have a maximum of 10 IPSec tunnels; this means that up to 10 networks can be associated with the same tunnel interface on the firewall.

The tunnel interface must belong to a security zone to apply policy and it must be assigned to a virtual router in order to use the existing routing infrastructure. Ensure that the tunnel interface and the physical interface are assigned to the same virtual router so that the firewall can perform a route lookup and determine the appropriate tunnel to use.

Typically, the Layer 3 interface that the tunnel interface is attached to belongs to an external zone, for example the untrust zone. While the tunnel interface can be in the same security zone as the physical interface, for added security and better visibility, you can create a separate zone for the tunnel interface. If you create a separate zone for the tunnel interface, say a VPN zone, you will need to create security policies to enable traffic to flow between the VPN zone and the trust zone.

 © Palo Alto Networks, Inc.

To route traffic between the sites, a tunnel interface does not require an IP address. An IP address is only required if you want to enable tunnel monitoring or if you are using a dynamic routing protocol to route traffic across the tunnel. With dynamic routing, the tunnel IP address serves as the next hop IP address for routing traffic to the VPN tunnel.

If you are configuring the Palo Alto Networks firewall with a VPN peer that performs policy-based VPN, you must configure a local and remote Proxy ID when setting up the IPSec tunnel. Each peer compares the Proxy-IDs configured on it with what is actually received in the packet in order to allow a successful IKE phase 2 negotiation. If multiple tunnels are required, configure unique Proxy IDs for each tunnel interface; a tunnel interface can have a maximum of 250 Proxy IDs. Each Proxy ID counts towards the IPSec VPN tunnel capacity of the firewall, and the tunnel capacity varies by the firewall model.

See Set Up an IPSec Tunnel for configuration details.

## Tunnel Monitoring

For a VPN tunnel, you can check connectivity to a destination IP address across the tunnel. The network monitoring profile on the firewall allows you to verify connectivity (using ICMP) to a destination IP address or a next hop at a specified polling interval, and to specify an action on failure to access the monitored IP address.

If the destination IP is unreachable, you either configure the firewall to wait for the tunnel to recover or configure automatic failover to another tunnel. In either case, the firewall generates a system log that alerts you to a tunnel failure and renegotiates the IPSec keys to accelerate recovery.

The default monitoring profile is configured to wait for the tunnel to recover; the polling interval is 3 seconds and the failure threshold is 5.

See Set Up Tunnel Monitoring for configuration details.

## Internet Key Exchange (IKE) for VPN

The IKE process allows the VPN peers at both ends of the tunnel to encrypt and decrypt packets using mutually agreed-upon keys or certificate and method of encryption. The IKE process occurs in two phases: IKE Phase 1 and IKE Phase 2. Each of these phases use keys and encryption algorithms that are defined using cryptographic profiles— IKE crypto profile and IPSec crypto profile—and the result of the IKE negotiation is a Security Association (SA). An SA is a set of mutually agreed-upon keys and algorithms that are used by both VPN peers to allow the flow of data across the VPN tunnel. The following illustration depicts the key exchange process for setting up the VPN tunnel:

© Palo Alto Networks, Inc.

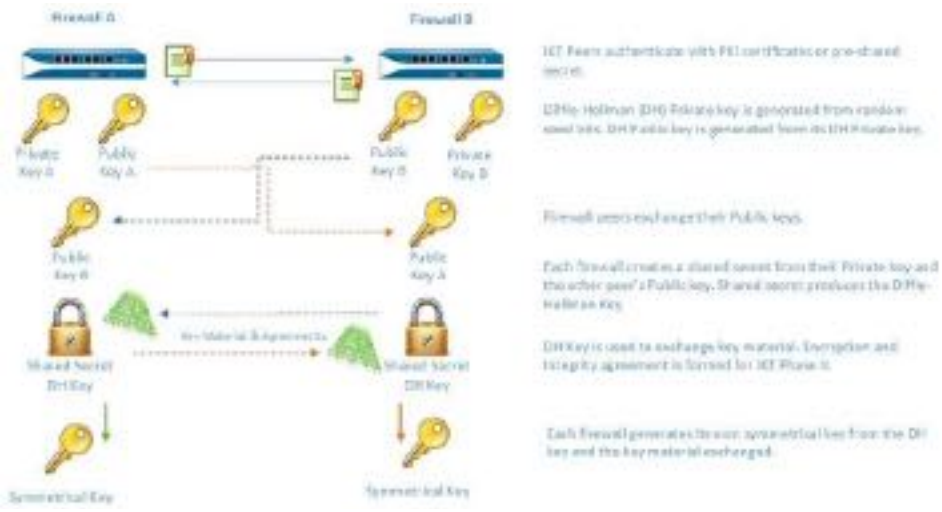

## IKE Phase 1

In this phase, the firewalls use the parameters defined in the IKE Gateway configuration and the IKE Crypto profile to authenticate each other and set up a secure control channel. IKE Phase supports the use of preshared keys or digital certificates (which use public key infrastructure, PKI) for mutual authentication of the VPN peers. Preshared keys are a simple solution for securing smaller networks because they do not require the support of a PKI infrastructure. Digital certificates can be more convenient for larger networks or implementations that require stronger authentication security.

When using certificates, make sure that the CA issuing the certificate is trusted by both gateway peers and that the maximum length of certificates in the certificate chain is 5 or less. With IKE fragmentation enabled, the firewall can reassemble IKE messages with up to 5 certificates in the certificate chain and successfully establish a VPN tunnel.

The IKE Crypto profile defines the following options that are used in the IKE SA negotiation:

- Diffie-Hellman (DH) group for generating symmetrical keys for IKE.

  The Diffie-Hellman algorithm uses the private key of one party and the public key of the other to create a shared secret, which is an encrypted key that both VPN tunnel peers share. The DH groups supported on the firewall are: Group 1—768 bits, Group 2—1024 bits (default), Group 5—1536 bits, Group 14—2048 bits, Group 19—256-bit elliptic curve group, and Group 20—384-bit elliptic curve group.

- Authentication algorithms—sha1, sha 256, sha 384, sha 512, or md5

- Encryption algorithms—3des, aes-128-cbc, aes-192-cbc, or aes-256-cbc

## IKE Phase 2

After the tunnel is secured and authenticated, in Phase 2 the channel is further secured for the transfer of data between the networks. IKE Phase 2 uses the keys that were established in Phase 1 of the process and the IPSec Crypto profile, which defines the IPSec protocols and keys used for the SA in IKE Phase 2.

 © Palo Alto Networks, Inc.

The IPSEC uses the following protocols to enable secure communication:

- Encapsulating Security Payload (ESP)—Allows you to encrypt the entire IP packet, and authenticate the source and verify integrity of the data. While ESP requires that you encrypt and authenticate the packet, you can choose to only encrypt or only authenticate by setting the encryption option to Null; using encryption without authentication is discouraged.
- Authentication Header (AH)—Authenticates the source of the packet and verifies data integrity. AH does not encrypt the data payload and is unsuited for deployments where data privacy is important. AH is commonly used when the main concern is to verify the legitimacy of the peer, and data privacy is not required.

**Table: Algorithms Supported for IPSEC Authentication and Encryption**

| ESP | AH |
|---|---|
| Diffie Hellman (DH) exchange options supported | |
| • Group 1—768 bits<br>• Group 2—1024 bits (the default)<br>• Group 5—1536 bits<br>• Group 14—2048 bits.<br>• Group 19— 256-bit elliptic curve group<br>• Group 20—384-bit elliptic curve group<br>• no-pfs—By default, perfect forward secrecy (PFS) is enabled, which means a new DH key is generated in IKE phase 2 using one of the groups listed above. This key is independent of the keys exchanged in IKE phase1 and provides better data transfer security. If you select no-pfs, the DH key created at phase 1 is not renewed and a single key is used for the IPSec SA negotiations. Both VPN peers must be enabled or disabled for PFS. | |
| Encryption algorithms supported | |
| • 3des | Triple Data Encryption Standard (3DES) with a security strength of 112 bits |
| • aes-128-cbc | Advanced Encryption Standard (AES) using cipher block chaining (CBC) with a security strength of 128 bits |
| • aes-192-cbc | AES using CBC with a security strength of 192 bits |
| • aes-256-cbc | AES using CBC with a security strength of 256 bits |
| • aes-128-ccm | AES using Counter with CBC-MAC (CCM) with a security strength of 128 bits |
| • aes-128-gcm | AES using Galois/Counter Mode (GCM) with a security strength of 128 bits |
| • aes-256-gcm | AES using GCM with a security strength of 256 bits |
| Authentication algorithms supported | |
| • md5 | • md5 |

| ESP | AH |
|---|---|
| • sha 1 | • sha 1 |
| • sha 256 | • sha 256 |
| • sha 384 | • sha 384 |
| • sha512 | • sha 512 |

## Methods of Securing IPSec VPN Tunnels (IKE Phase 2)

IPSec VPN tunnels can be secured using manual keys or auto keys. In addition, IPSec configuration options include Diffie-Hellman Group for key agreement, and/or an encryption algorithm and a hash for message authentication.

- **Manual Key**—Manual key is typically used if the Palo Alto Networks firewall is establishing a VPN tunnel with a legacy device, or if you want to reduce the overhead of generating session keys. If using manual keys, the same key must be configured on both peers.

  Manual keys are not recommended for establishing a VPN tunnel because the session keys can be compromised when relaying the key information between the peers; if the keys are compromised, the data transfer is no longer secure.

- **Auto Key**— Auto Key allows you to automatically generate keys for setting up and maintaining the IPSec tunnel based on the algorithms defined in the IPSec Crypto profile.

# IKEv2

An IPSec VPN gateway uses IKEv1 or IKEv2 to negotiate the IKE security association (SA) and IPSec tunnel. IKEv2 is defined in RFC 5996.

Unlike IKEv1, which uses Phase 1 SA and Phase 2 SA, IKEv2 uses a child SA for Encapsulating Security Payload (ESP) or Authentication Header (AH), which is set up with an IKE SA.

NAT traversal (NAT-T) must be enabled on both gateways if you have NAT occurring on a device that sits between the two gateways. A gateway can see only the public (globally routable) IP address of the NAT device.

IKEv2 provides the following benefits over IKEv1:

- Tunnel endpoints exchange fewer messages to establish a tunnel. IKEv2 uses four messages; IKEv1 uses either nine messages (in main mode) or six messages (in aggressive mode).
- Built-in NAT-T functionality improves compatibility between vendors.
- Built-in health check automatically re-establishes a tunnel if it goes down. The liveness check replaces the Dead Peer Detection used in IKEv1.
- Supports traffic selectors (one per exchange). The traffic selectors are used in IKE negotiations to control what traffic can access the tunnel.

 © Palo Alto Networks, Inc.

- Supports Hash and URL certificate exchange to reduce fragmentation.
- Resiliency against DoS attacks with improved peer validation. An excessive number of half-open SAs can trigger cookie validation.

Before configuring IKEv2, you should be familiar with the following concepts:

- Liveness Check
- Cookie Activation Threshold and Strict Cookie Validation
- Traffic Selectors
- Hash and URL Certificate Exchange
- SA Key Lifetime and Re-Authentication Interval

After you Set Up an IKE Gateway, if you chose IKEv2, perform the following optional tasks related to IKEv2 as required by your environment:

- Export a Certificate for a Peer to Access Using Hash and URL
- Import a Certificate for IKEv2 Gateway Authentication
- Change the Key Lifetime or Authentication Interval for IKEv2
- Change the Cookie Activation Threshold for IKEv2
- Configure IKEv2 Traffic Selectors

## Liveness Check

The liveness check for IKEv2 is similar to Dead Peer Detection (DPD), which IKEv1 uses as the way to determine whether a peer is still available.

In IKEv2, the liveness check is achieved by any IKEv2 packet transmission or an empty informational message that the gateway sends to the peer at a configurable interval, five seconds by default. If necessary, the sender attempts the retransmission up to ten times. If it doesn't get a response, the sender closes and deletes the IKE_SA and corresponding CHILD_SAs. The sender will start over by sending out another IKE_SA_INIT message.

## Cookie Activation Threshold and Strict Cookie Validation

Cookie validation is always enabled for IKEv2; it helps protect against half-SA DoS attacks. You can configure the global threshold number of half-open SAs that will trigger cookie validation. You can also configure individual IKE gateways to enforce cookie validation for every new IKEv2 SA.

- The **Cookie Activation Threshold** is a global VPN session setting that limits the number of simultaneous half-opened IKE SAs (default is 500). When the number of half-opened IKE SAs exceeds the **Cookie Activation Threshold**, the Responder will request a cookie, and the Initiator must respond with an IKE_SA_INIT containing a cookie to validate the connection. If the cookie validation is successful, another SA can be initiated. A value of 0 means that cookie validation is always on.

The Responder does not maintain a state of the Initiator, nor does it perform a Diffie-Hellman key exchange, until the Initiator returns the cookie. IKEv2 cookie validation mitigates a DoS attack that would try to leave numerous connections half open.

The **Cookie Activation Threshold** must be lower than the **Maximum Half Opened SA** setting. If you Change the Cookie Activation Threshold for IKEv2 to a very high number (for example, 65534) and the **Maximum Half Opened SA** setting remained at the default value of 65535, cookie validation is essentially disabled.

- You can enable **Strict Cookie Validation** if you want cookie validation performed for every new IKEv2 SA a gateway receives, regardless of the global threshold. **Strict Cookie Validation** affects only the IKE gateway being configured and is disabled by default. If **Strict Cookie Validation** is disabled, the system uses the **Cookie Activation Threshold** to determine whether a cookie is needed or not.

## Traffic Selectors

In IKEv1, a firewall that has a route-based VPN needs to use a local and remote Proxy ID in order to set up an IPSec tunnel. Each peer compares its Proxy IDs with what it received in the packet in order to successfully negotiate IKE Phase 2. IKE Phase 2 is about negotiating the SAs to set up an IPSec tunnel. (For more information on Proxy IDs, see Tunnel Interface.)

In IKEv2, you can Configure IKEv2 Traffic Selectors, which are components of network traffic that are used during IKE negotiation. Traffic selectors are used during the CHILD_SA (tunnel creation) Phase 2 to set up the tunnel and to determine what traffic is allowed through the tunnel. The two IKE gateway peers must negotiate and agree on their traffic selectors; otherwise, one side narrows its address range to reach agreement. One IKE connection can have multiple tunnels; for example, you can assign different tunnels to each department to isolate their traffic. Separation of traffic also allows features such as QoS to be implemented.

The IPv4 and IPv6 traffic selectors are:

- **Source IP address**—A network prefix, address range, specific host, or wildcard.
- **Destination IP address**—A network prefix, address range, specific host, or wildcard.
- **Protocol**—A transport protocol, such as TCP or UDP.
- **Source port**—The port where the packet originated.
- **Destination port**—The port the packet is destined for.

During IKE negotiation, there can be multiple traffic selectors for different networks and protocols. For example, the Initiator might indicate that it wants to send TCP packets from 172.168.0.0/16 through the tunnel to its peer, destined for 198.5.0.0/16. It also wants to send UDP packets from 172.17.0.0/16 through the same tunnel to the same gateway, destined for 0.0.0.0 (any network). The peer gateway must agree to these traffic selectors so that it knows what to expect.

It is possible that one gateway will start negotiation using a traffic selector that is a more specific IP address than the IP address of the other gateway.

 © Palo Alto Networks, Inc.

- For example, gateway A offers a source IP address of 172.16.0.0/16 and a destination IP address of 192.16.0.0/16. But gateway B is configured with 0.0.0.0 (any source) as the source IP address and 0.0.0.0 (any destination) as the destination IP address. Therefore, gateway B narrows down its source IP address to 192.16.0.0/16 and its destination address to 172.16.0.0/16. Thus, the narrowing down accommodates the addresses of gateway A and the traffic selectors of the two gateways are in agreement.
- If gateway B (configured with source IP address 0.0.0.0) is the Initiator instead of the Responder, gateway A will respond with its more specific IP addresses, and gateway B will narrow down its addresses to reach agreement.

## Hash and URL Certificate Exchange

IKEv2 supports Hash and URL Certificate Exchange, which is used during an IKEv2 negotiation of an SA. You store the certificate on an HTTP server, which is specified by a URL. The peer fetches the certificate from the server based on receiving the URL to the server. The hash is used to check whether the content of the certificate is valid or not. Thus, the two peers exchange certificates with the HTTP CA rather than with each other.

The hash part of Hash and URL reduces the message size and thus Hash and URL is a way to reduce the likelihood of packet fragmentation during IKE negotiation. The peer receives the certificate and hash that it expects, and thus IKE Phase 1 has validated the peer. Reducing fragmentation occurrences helps protect against DoS attacks.

You can enable the Hash and URL certificate exchange when configuring an IKE gateway by selecting **HTTP Certificate Exchange** and entering the **Certificate URL**. The peer must also use Hash and URL certificate exchange in order for the exchange to be successful. If the peer cannot use Hash and URL, X.509 certificates are exchanged similarly to how they are exchanged in IKEv1.

If you enable the Hash and URL certificate exchange, you must export your certificate to the certificate server if it is not already there. When you export the certificate, the file format should be **Binary Encoded Certificate (DER)**. See Export a Certificate for a Peer to Access Using Hash and URL.

## SA Key Lifetime and Re-Authentication Interval

In IKEv2, two IKE crypto profile values, **Key Lifetime** and **IKEv2 Authentication Multiple**, control the establishment of IKEv2 IKE SAs. The key lifetime is the length of time that a negotiated IKE SA key is effective. Before the key lifetime expires, the SA must be re-keyed; otherwise, upon expiration, the SA must begin a new IKEv2 IKE SA re-key. The default value is 8 hours.

The re-authentication interval is derived by multiplying the **Key Lifetime** by the **IKEv2 Authentication Multiple**. The authentication multiple defaults to 0, which disables the re-authentication feature.

The range of the authentication multiple is 0-50. So, if you were to configure an authentication multiple of 20, for example, the system would perform re-authentication every 20 re-keys, which is every 160 hours. That means the gateway could perform Child SA creation for 160 hours before the gateway must re-authenticate with IKE to recreate the IKE SA from scratch.

In IKEv2, the Initiator and Responder gateways have their own key lifetime value, and the gateway with the shorter key lifetime is the one that will request that the SA be re-keyed.

© Palo Alto Networks, Inc.

# Set Up Site-to-Site VPN

To set up site-to-site VPN:

- Make sure that your Ethernet interfaces, virtual routers, and zones are configured properly. For more information, see Configure Interfaces and Zones.
- Create your tunnel interfaces. Ideally, put the tunnel interfaces in a separate zone, so that tunneled traffic can use different policies.
- Set up static routes or assign routing protocols to redirect traffic to the VPN tunnels. To support dynamic routing (OSPF, BGP, RIP are supported), you must assign an IP address to the tunnel interface.
- Define IKE gateways for establishing communication between the peers across each end of the VPN tunnel; also define the cryptographic profile that specifies the protocols and algorithms for identification, authentication, and encryption to be used for setting up VPN tunnels in IKEv1 Phase 1. See Set Up an IKE Gateway and Define IKE Crypto Profiles.
- Configure the parameters that are needed to establish the IPSec connection for transfer of data across the VPN tunnel; See Set Up an IPSec Tunnel. For IKEv1 Phase-2, see Define IPSec Crypto Profiles.
- (Optional) Specify how the firewall will monitor the IPSec tunnels. See Set Up Tunnel Monitoring.
- Define security policies to filter and inspect the traffic.

If there is a deny rule at the end of the security rulebase, intra-zone traffic is blocked unless otherwise allowed. Rules to allow IKE and IPSec applications must be explicitly included above the deny rule.

When these tasks are complete, the tunnel is ready for use. Traffic destined for the zones/addresses defined in policy is automatically routed properly based on the destination route in the routing table, and handled as VPN traffic. For a few examples on site-to-site VPN, see Site-to-Site VPN Quick Configs.

For troubleshooting purposes, you can Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel.

 © Palo Alto Networks, Inc.

## Set Up an IKE Gateway

To set up a VPN tunnel, the VPN peers or gateways must authenticate each other using preshared keys or digital certificates and establish a secure channel in which to negotiate the IPSec security association (SA) that will be used to secure traffic between the hosts on each side.

| Set Up an IKE Gateway | | |
|---|---|---|
| Step 1 | Define the IKE Gateway. | 1. Select **Network > Network Profiles > IKE Gateways**, click **Add**, and on the **General** tab, enter the **Name** of the gateway.<br>2. For **Version**, select **IKEv1 only mode**, **IKEv2 only mode**, or **IKEv2 preferred mode**. The IKE gateway begins its negotiation with its peer in the mode specified here. If you select **IKEv2 preferred mode**, the two peers will use IKEv2 if the remote peer supports it; otherwise they will use IKEv1.<br>(The **Version** selection also determines which options are available on the **Advanced Options** tab.) |
| Step 2 | Establish the local endpoint of the tunnel (gateway). | 1. For **Address Type**, click **IPv4** or **IPv6**.<br>2. Select the physical, outgoing **Interface** on the firewall where the local gateway resides.<br>3. From the **Local IP Address** drop-down, select the IP address that will be used as the endpoint for the VPN connection. This is the external-facing interface with a publicly routable IP address on the firewall. |
| Step 3 | Establish the peer at the far end of the tunnel (gateway). | 1. Select the **Peer IP Type** to be a **Static** or **Dynamic** address assignment.<br>2. If the **Peer IP Address** is static, enter the IP address of the peer. |
| Step 4 | Specify how the peer is authenticated. | Select the **Authentication** method: **Pre-Shared Key** or **Certificate**. If you choose Pre-Shared Key, proceed to Step 5. If you choose Certificate, proceed to Step 6. |

| Set Up an IKE Gateway | |
|---|---|
| Step 5 Configure a pre-shared key. | 1. Enter a **Pre-shared Key**, which is the security key to use for authentication across the tunnel. Re-enter the value to **Confirm Pre-shared Key**.<br> Generate a key that is difficult to crack with dictionary attacks; use a pre-shared key generator, if necessary.<br>2. For **Local Identification**, choose from the following types and enter a value that you determine: **FQDN (hostname)**, **IP address**, **KEYID (binary format ID string in HEX)**, **User FQDN (email address)**. Local identification defines the format and identification of the local gateway. If no value is specified, the local IP address will be used as the local identification value.<br>3. For **Peer Identification**, choose from the following types and enter the value: **FQDN (hostname)**, **IP address**, **KEYID (binary format ID string in HEX)**, **User FQDN (email address)**. Peer identification defines the format and identification of the peer gateway. If no value is specified, the peer IP address will be used as the peer identification value.<br>4. Proceed to Step 7 and continue from there. |

 © Palo Alto Networks, Inc.

| Set Up an IKE Gateway | |
|---|---|
| Step 6 Configure certificate-based authentication. Perform the remaining steps in this procedure if you selected **Certificate** as the method of authenticating the peer gateway at the opposite end of the tunnel. | 1. Select a **Local Certificate** that is already on the firewall from the drop-down, or **Import** a certificate, or **Generate** to create a new certificate.<br>• If you want to **Import** a certificate, Import a Certificate for IKEv2 Gateway Authentication and then return to this task.<br>• If you want to **Generate** a new certificate, Generate a Certificate on the Device and then return to this task.<br>2. Click the **HTTP Certificate Exchange** check box if you want to configure Hash and URL (IKEv2 only). For an HTTP certificate exchange, enter the **Certificate URL**. For more information, see Hash and URL Certificate Exchange.<br>3. Select the **Local Identification** type from the following: **Distinguished Name (Subject)**, **FQDN (hostname)**, **IP address**, **User FQDN (email address)**, and enter the value. Local identification defines the format and identification of the local gateway.<br>4. Select the **Peer Identification** type from the following: **Distinguished Name (Subject)**, **FQDN (hostname)**, **IP address**, **User FQDN (email address)**, and enter the value. Peer identification defines the format and identification of the peer gateway.<br>5. Select one type of **Peer ID Check**:<br>• **Exact**—Check this to ensure that the local setting and peer IKE ID payload match exactly.<br>• **Wildcard**—Check this to allow the peer identification to match as long as every character before the wildcard (*) matches. The characters after the wildcard need not match.<br>6. Click **Permit peer identification and certificate payload identification mismatch** if you want to allow a successful IKE SA even when the peer identification does not match the peer identification in the certificate.<br>7. Choose a **Certificate Profile** from the drop-down. A certificate profile contains information about how to authenticate the peer gateway.<br>8. Click **Enable strict validation of peer's extended key use** if you want to strictly control how the key can be used. |

© Palo Alto Networks, Inc.

| Set Up an IKE Gateway | |
|---|---|
| Step 7 Configure advanced options for the gateway. | 1. Select the **Advanced Options** tab.<br>2. In the Common Options section, **Enable Passive Mode** if you want the firewall to only respond to IKE connection requests and never initiate them.<br>3. **Enable NAT Traversal** if you have a device performing NAT between the gateways, to have UDP encapsulation used on IKE and UDP protocols, enabling them to pass through intermediate NAT devices.<br>4. If you chose **IKEv1 only mode** earlier, on the IKEv1 tab:<br>• Choose **auto**, **aggressive**, or **main** for the **Exchange Mode**. When a device is set to use **auto** exchange mode, it can accept both **main** mode and **aggressive** mode negotiation requests; however, whenever possible, it initiates negotiation and allows exchanges in **main** mode.<br>If the exchange mode is not set to **auto**, you must configure both peers with the same exchange mode to allow each peer to accept negotiation requests.<br>• Select an existing profile or keep the default profile from **IKE Crypto Profile** drop-down. For details on defining an IKE Crypto profile, see Define IKE Crypto Profiles.<br>• (Only if using certificate-based authentication and the exchange mode is not set to **aggressive** mode) Click **Enable Fragmentation** to enable the firewall to operate with IKE Fragmentation.<br>• Click **Dead Peer Detection** and enter an **Interval** (range is 2-100 seconds). For **Retry,** define the time to delay (range is 2-100 seconds) before attempting to re-check availability. Dead peer detection identifies inactive or unavailable IKE peers by sending an IKE phase 1 notification payload to the peer and waiting for an acknowledgment.<br>5. If you chose **IKEv2 only mode** or **IKEv2 preferred mode** in Step 1, on the IKEv2 tab:<br>• Select an **IKE Crypto Profile** from the drop-down, which configures IKE Phase 1 options such as the DH group, hash algorithm, and ESP authentication. For information about IKE crypto profiles, see IKE Phase 1.<br>• Enable **Strict Cookie Validation** if you want to always enforce cookie validation on IKEv2 SAs for this gateway. See Cookie Activation Threshold and Strict Cookie Validation.<br>• **Enable Liveness Check** and enter an **Interval (sec)** (default is 5) if you want to have the gateway send a message request to its gateway peer, requesting a response. If necessary, the Initiator attempts the liveness check up to 10 times. If it doesn't get a response, the Initiator closes and deletes the IKE_SA and CHILD_SA. The Initiator will start over by sending out another IKE_SA_INIT. |


© Palo Alto Networks, Inc.

| Set Up an IKE Gateway | |
|---|---|
| Step 8 Save the changes. | Click **OK** and **Commit**. |

## Export a Certificate for a Peer to Access Using Hash and URL

IKEv2 supports Hash and URL Certificate Exchange as a method of having the peer at the remote end of the tunnel fetch the certificate from a server where you have exported the certificate. Perform this task to export your certificate to that server.

❒ You must have already created a certificate using **Device > Certificate Management**.

| Export a Certificate for Hash and URL | |
|---|---|
| Export a certificate for a peer to access using Hash and URL certificate exchange. | 1. Select **Device > Certificates**, and if your platform supports multiple virtual systems, for **Location**, select the appropriate virtual system.<br>2. On the **Device Certificates** tab, select the certificate to **Export** to the server.<br>The status of the certificate should be valid, not expired. The firewall will not stop you from exporting an invalid certificate.<br>3. For **File Format**, select **Binary Encoded Certificate (DER)**.<br>4. Leave **Export private key** clear. Exporting the private key is unnecessary for Hash and URL.<br>5. Click **OK**. |

## Import a Certificate for IKEv2 Gateway Authentication

Perform this task if you are authenticating a peer for an IKEv2 gateway and you did not use a local certificate already on the firewall; you want to import a certificate from elsewhere.

This task presumes that you selected **Network > IKE Gateways**, added a gateway, and for **Local Certificate**, you clicked **Import**.

| Import a Certificate for IKEv2 Gateway Authentication | |
|---|---|
| Step 1 Import a certificate. | 1. Select **Network > IKE Gateways**, **Add** a gateway, and on the **General** tab, for **Authentication**, select **Certificate**. For **Local Certificate**, click **Import**.<br>2. In the Import Certificate window, enter a **Certificate Name** for the certificate you are importing.<br>3. Select **Shared** if this certificate is to be shared among multiple virtual systems.<br>4. For **Certificate File**, **Browse** to the certificate file. Click on the file name and click **Open**, which populates the **Certificate File** field.<br>5. For **File Format**, select one of the following:<br>• **Base64 Encoded Certificate (PEM)**—Contains the certificate, but not the key. It is cleartext.<br>• **Encrypted Private Key and Certificate (PKCS12)**—Contains both the certificate and the key.<br>6. Select **Import private key** if the key is in a different file from the certificate file. The key is optional, with the following exception:<br>• You must import a key if you set the **File Format** to **PEM**. Enter a **Key file** by clicking **Browse** and navigating to the key file to import.<br>• Enter a **Passphrase** and **Confirm Passphrase**.<br>7. Click **OK**. |
| Step 2 After you perform this task, return to Configure an IKEv2 Gateway and resume Step 6. | |

## Change the Key Lifetime or Authentication Interval for IKEv2

This task is optional; the default setting of the IKEv2 IKE SA re-key lifetime is 8 hours. The default setting of the IKEv2 Authentication Multiple is 0, meaning the re-authentication feature is disabled. For more information, see SA Key Lifetime and Re-Authentication Interval.

To change the default values, perform the following task. A prerequisite is that an IKE crypto profile already exists.

 © Palo Alto Networks, Inc.

| Change the SA Key Lifetime or Authentication Interval | | |
|---|---|---|
| Step 1 | Change the SA key lifetime or authentication interval for an IKE Crypto profile. | 1. Select **Network > Network Profiles > IKE Crypto** and select the IKE Crypto profile that applies to the local gateway.<br>2. For the **Key Lifetime**, select a unit (**Seconds**, **Minutes**, **Hours**, or **Days**) and enter a value. The minimum is three minutes.<br>3. For **IKE Authentication Multiple**, enter a value, which is multiplied by the lifetime to determine the re-authentication interval. |
| Step 2 | Save the configuration. | Click **OK** and **Commit**. |

## Change the Cookie Activation Threshold for IKEv2

Perform the following task if you want a firewall to have a threshold different from the default setting of 500 half-opened SA sessions before cookie validation is required. For more information about cookie validation, see Cookie Activation Threshold and Strict Cookie Validation.

| Change the Cookie Activation Threshold | | |
|---|---|---|
| Step 1 | Change the Cookie Activation Threshold. | 1. Select **Device > Setup > Session** and edit the VPN Session Settings. For **Cookie Activation Threshold**, enter the maximum number of half-opened SAs that are allowed before the responder requests a cookie from the initiator (range is 0-65535; default: is 500).<br>2. Click **OK**. |
| Step 2 | Save the configuration | Click **OK** and **Commit**. |

## Configure IKEv2 Traffic Selectors

| Configure Traffic Selectors for IKEv2 | | |
|---|---|---|
| Step 1 | Configure Traffic Selectors. | 1. Select **Network > IPSec Tunnels > Proxy IDs**.<br>2. Select the **IPv4** or **IPv6** tab.<br>3. Click **Add** and enter the **Name** in the **Proxy ID** field.<br>4. In the **Local** field, enter the **Source IP Address**.<br>5. In the **Remote** field, enter the **Destination IP Address**.<br>6. In the **Protocol** field, select the transport protocol (**TCP** or **UDP**) from the drop-down.<br>7. Click **OK**. |

# Define Cryptographic Profiles

A cryptographic profile specifies the ciphers used for authentication and/or encryption between two IKE peers, and the lifetime of the key. The time period between each renegotiation is known as the lifetime; when the specified time expires, the firewall renegotiates a new set of keys.

For securing communication across the VPN tunnel, the firewall requires IKE and IPSec cryptographic profiles for completing IKE phase 1 and phase 2 negotiations, respectively. The firewall includes a default IKE crypto profile and a default IPSec crypto profile that is ready for use.

- Define IKE Crypto Profiles
- Define IPSec Crypto Profiles

## Define IKE Crypto Profiles

The IKE crypto profile is used to set up the encryption and authentication algorithms used for the key exchange process in IKE Phase 1, and lifetime of the keys, which specifies how long the keys are valid. To invoke the profile, you must attach it to the IKE Gateway configuration.

All IKE gateways configured on the same interface or local IP address must use the same crypto profile.

| **Define an IKE Crypto Profile** | | |
|---|---|---|
| Step 1 | Create a new IKE profile. | 1. Select **Network > Network Profiles > IKE Crypto** and select **Add**.<br>2. Enter a **Name** for the new profile. |
| Step 2 | Specify the DH Group (Diffie–Hellman group) for key exchange, and the Authentication and Encryption algorithms. | Click **Add** in the corresponding sections (DH Group, Authentication, and Encryption) and select from the drop-downs.<br>If you are not certain of what the VPN peers support, add multiple groups or algorithms in the order of most-to-least secure as follows; the peers negotiate the strongest supported group or algorithm to establish the tunnel:<br>• DH Group—**group20**, **group19**, **group14**, **group5**, **group2**, and **group1**.<br>• Authentication—**sha512**, **sha384**, **sha256**, **sha1**, **md5**.<br>• Encryption—**aes-256-cbc**, **aes-192-cbc**, **aes-128-cbc**, **3des**. |
| Step 3 | Specify the duration for which the key is valid and the re-authentication interval.<br>For details, see SA Key Lifetime and Re-Authentication Interval. | 1. In the **Key Lifetime** fields, specify the period (in seconds, minutes, hours, or days) for which the key is valid. (Range is 3 minutes to 365 days; default is 8 hours.) When the key expires, the firewall renegotiates a new key. A lifetime is the period between each renegotiation.<br>2. For the **IKEv2 Authentication Multiple**, specify a value (range is 0-50) that is multiplied by the **Key Lifetime** to determine the authentication count. The default value of 0 disables the re-authentication feature. |

 © Palo Alto Networks, Inc.

| Define an IKE Crypto Profile | | |
|---|---|---|
| Step 4 | Save your IKE Crypto profile. | Click **OK** and click **Commit**. |
| Step 5 | Attach the IKE Crypto profile to the IKE Gateway configuration. | See Step 7 in Set Up an IKE Gateway. |

## Define IPSec Crypto Profiles

The IPSec crypto profile is invoked in IKE Phase 2. It specifies how the data is secured within the tunnel when Auto Key IKE is used to automatically generate keys for the IKE SAs.

| Define the IPSec Crypto Profile | | |
|---|---|---|
| Step 1 | Create a new IPSec profile. | 1. Select **Network > Network Profiles > IPSec Crypto** and select **Add**.<br>2. Enter a **Name** for the new profile.<br>3. Select the **IPSec Protocol**—ESP or AH—that you want to apply to secure the data as it traverses across the tunnel.<br>4. Click **Add** and select the **Authentication** and **Encryption** algorithms for ESP, and **Authentication** algorithms for AH, so that the IKE peers can negotiate the keys for the secure transfer of data across the tunnel.<br>If you are not certain of what the IKE peers support, add multiple algorithms in the order of most-to-least secure as follows; the peers negotiate the strongest supported algorithm to establish the tunnel:<br>• Encryption—**aes-256-gcm**, **aes-256-cbc**, **aes-192-cbc**, **aes-128-gcm**, **aes-128-ccm** (the VM-Series firewall doesn't support this option), **aes-128-cbc**, **3des**.<br>• Authentication—**sha512**, **sha384**, **sha256**, **sha1**, **md5**. |
| Step 2 | Select the DH Group to use for the IPSec SA negotiations in IKE phase 2. | Select the key strength that you want to use from the **DH Group** drop-down.<br>If you are not certain of what the VPN peers support, add multiple groups in the order of most-to-least secure as follows; the peers negotiate the strongest supported group to establish the tunnel: **group20**, **group19**, **group14**, **group5**, **group2**, and **group1**.<br>Select **no-pfs** if you do not want to renew the key that was created at phase 1; the current key is reused for the IPSEC SA negotiations. |

| Define the IPSec Crypto Profile | | |
|---|---|---|
| Step 3 | Specify the duration of the key—time and volume of traffic. | Using a combination of time and traffic volume allows you to ensure safety of data.<br>Select the **Lifetime** or time period for which the key is valid in seconds, minutes, hours, or days (range is 3 minutes to 365 days). When the specified time expires, the firewall will renegotiate a new set of keys.<br>Select the **Lifesize** or volume of data after which the keys must be renegotiated. |
| Step 4 | Save your IPSec profile. | Click **OK** and click **Commit**. |
| Step 5 | Attach the IPSec Profile to an IPSec tunnel configuration. | See Step 4 in Set Up an IPSec Tunnel. |


© Palo Alto Networks, Inc.

## Set Up an IPSec Tunnel

The IPSec tunnel configuration allows you to authenticate and/or encrypt the data (IP packet) as it traverses across the tunnel.

If you are setting up the Palo Alto Networks firewall to work with a peer that supports policy-based VPN, you must define Proxy IDs. Devices that support policy-based VPN use specific security rules/policies or access-lists (source addresses, destination addresses and ports) for permitting interesting traffic through an IPSec tunnel. These rules are referenced during quick mode/IKE phase 2 negotiation, and are exchanged as Proxy-IDs in the first or the second message of the process. So, if you are configuring the Palo Alto Networks firewall to work with a policy-based VPN peer, for a successful phase 2 negotiation you must define the Proxy-ID so that the setting on both peers is identical. If the Proxy-ID is not configured, because the Palo Alto Networks firewall supports route-based VPN, the default values used as Proxy-ID are source ip: 0.0.0.0/0, destination ip: 0.0.0.0/0 and application: any; and when these values are exchanged with the peer, it results in a failure to set up the VPN connection.

**Set Up an IPSec Tunnel**

Step 1 Select **Network > IPSec Tunnels> General** and enter a **Name** for the new tunnel.

Step 2 Select the **Tunnel interface** that will be used to set up the IPSec tunnel.

To create a new tunnel interface:

1. Select **Network > Interfaces > Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down. Associating the tunnel interface with the same zone (and virtual router) as the external-facing interface on which the packets enter the firewall, mitigates the need to create inter-zone routing.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example vpn-corp), and click **OK**.
4. In the **Virtual Router** drop-down, select **default**.
5. (Optional) If you want to assign an IPv4 address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.31.32.1/32.
6. If you want to assign an IPv6 address to the tunnel interface, see Step 3.
7. To save the interface configuration, click **OK**.

| Set Up an IPSec Tunnel | |
|---|---|
| Step 3 (Optional) Enable IPv6 on the tunnel interface. | 1. Select the IPv6 tab on **Network > Interfaces > Tunnel > IPv6**.<br>2. Select the check box to **Enable IPv6 on the interface**.<br>This option allows you to route IPv6 traffic over an IPv4 IPSec tunnel and will provide confidentiality between IPv6 networks. The IPv6 traffic is encapsulated by IPv4 and then ESP. To route IPv6 traffic to the tunnel, you can use a static route to the tunnel, or use OSPFv3, or use a Policy-Based Forwarding (PBF) rule to direct traffic to the tunnel.<br>3. Enter the 64-bit extended unique **Interface ID** in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. By default, the firewall will use the EUI-64 generated from the physical interface's MAC address.<br>4. To enter an IPv6 **Address**, click **Add** and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. If Prefix is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.<br>a. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address.<br>b. Select **Anycast** to include routing through the nearest node. |
| Step 4 Select the type of key that will be used to secure the IPSec tunnel. | Continue to one of the following steps, depending on what type of key exchange you are using:<br>• Set up Auto Key exchange.<br>• Set up a Manual Key exchange. |
| • Set up Auto Key exchange. | 1. Select the IKE Gateway. To set up an IKE gateway, see Set Up an IKE Gateway.<br>2. (Optional) Select the default IPSec Crypto Profile. To create a new IPSec Profile, see Define IPSec Crypto Profiles. |


© Palo Alto Networks, Inc.

| Set Up an IPSec Tunnel | |
|---|---|
| • Set up a Manual Key exchange. | 1. Set up the parameters for the local firewall:<br>a. Specify the **SPI** for the local firewall. SPI is a 32-bit hexadecimal index that is added to the header for IPSec tunneling to assist in differentiating between IPSec traffic flows; it is used to create the SA required for establishing a VPN tunnel.<br>b. Select the **Interface** that will be the tunnel endpoint, and optionally select the IP address for the local interface that is the endpoint of the tunnel.<br>c. Select the protocol to be used—**AH** or **ESP**.<br>d. For AH, select the **Authentication** method from the drop-down and enter a **Key** and then **Confirm Key**.<br>e. For ESP, select the **Authentication** method from the drop-down and enter a **Key** and then **Confirm Key**. Then, select the **Encryption** method and enter a **Key** and then **Confirm Key**, if needed.<br>2. Set up the parameters that pertain to the remote VPN peer.<br>a. Specify the **SPI** for the remote peer.<br>b. Enter the **Remote Address**, the IP address of the remote peer. |
| Step 5 Protect against a replay attack.<br>A replay attack occurs when a packet is maliciously intercepted and retransmitted by the interceptor. | Select the **Show Advanced Options** check box, select **Enable Replay Protection** to detect and neutralize against replay attacks. |
| Step 6 Preserve the Type of Service header for the priority or treatment of IP packets. | In the **Show Advanced Options** section, select **Copy TOS Header**. This copies the Type of Service (TOS) header from the inner IP header to the outer IP header of the encapsulated packets in order to preserve the original TOS information. |
| Step 7 Enable Tunnel Monitoring.<br>You need to assign an IP address to the tunnel interface for monitoring. | To alert the device administrator to tunnel failures and to provide automatic failover to another tunnel interface:<br>1. Specify a **Destination IP** address on the other side of the tunnel to determine if the tunnel is working properly.<br>2. Select a **Profile** to determine the action on tunnel failure. To create a new profile, see Define a Tunnel Monitoring Profile. |

| Set Up an IPSec Tunnel | | |
|---|---|---|
| Step 8 | (Required only if the VPN peer uses policy-based VPN). Create a Proxy ID to identify the VPN peers. | 1. Select **Network > IPSec Tunnels** and click **Add**.<br>2. Select the **Proxy IDs** tab.<br>3. Select the **IPv4** or **IPv6** tab.<br>4. Click **Add** and enter the **Proxy ID** name.<br>5. Enter the **Local** IP address or subnet for the VPN gateway.<br>6. Enter the **Remote** address for the VPN gateway.<br>7. Select the **Protocol** from the drop-down:<br>• **Number**—Specify the protocol number (used for interoperability with third-party devices).<br>• **Any**—Allows TCP and/or UDP traffic.<br>• **TCP**—Specify the Local Port and Remote Port numbers.<br>• **UDP**—Specify the Local Port and Remote Port numbers.<br>8. Click **OK**. |
| Step 9 | Save your changes. | Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# Set Up Tunnel Monitoring

To provide uninterrupted VPN service, you can use the Dead Peer Detection capability along with the tunnel monitoring capability on the firewall. You can also monitor the status of the tunnel. These monitoring tasks are described in the following sections:

- Define a Tunnel Monitoring Profile
- View the Status of the Tunnels

## Define a Tunnel Monitoring Profile

A tunnel monitoring profile allows you to verify connectivity between the VPN peers; you can configure the tunnel interface to ping a destination IP address at a specified interval and specify the action if the communication across the tunnel is broken.

| Define a Tunnel Monitoring Profile |
|---|
| 1. Select **Network > Network Profiles > Monitor**. A default tunnel monitoring profile is available for use. |
| 2. Click **Add**, and enter a **Name** for the profile. |
| 3. Select the **Action** if the destination IP address is unreachable.<br>• **Wait Recover**—the firewall waits for the tunnel to recover. It continues to use the tunnel interface in routing decisions as if the tunnel were still active.<br>• **Fail Over**—forces traffic to a back-up path if one is available. The firewall disables the tunnel interface, and thereby disables any routes in the routing table that use the interface.<br>In either case, the firewall attempts to accelerate the recovery by negotiating new IPSec keys. |
| 4. Specify the **Interval** and **Threshold** to trigger the specified action.<br>The threshold specifies the number of heartbeats to wait before taking the specified action. The range is 2-100 and the default is 5.<br>The Interval measures the time between heartbeats. The range is 2-10 and the default is 3 seconds. |
| 5. Attach the monitoring profile to the IPsec Tunnel configuration. See Enable Tunnel Monitoring. |

## View the Status of the Tunnels

The status of the tunnel informs you about whether or not valid IKE phase-1 and phase-2 SAs have been established, and whether the tunnel interface is up and available for passing traffic.

Because the tunnel interface is a logical interface, it cannot indicate a physical link status. Therefore, you must enable tunnel monitoring so that the tunnel interface can verify connectivity to an IP address and determine if the path is still usable. If the IP address is unreachable, the firewall will either wait for the tunnel to recover or failover. When a failover occurs, the existing tunnel is torn down and routing changes are triggered to set up a new tunnel and redirect traffic.

© Palo Alto Networks, Inc.

| View Tunnel Status |
|---|
| 1. Select **Network > IPSec Tunnels**. |
| 2. View the **Tunnel Status**.<br>• Green indicates a valid IPSec SA tunnel.<br>• Red indicates that IPSec SA is not available or has expired. |
| 3. View the **IKE Gateway Status**.<br>• Green indicates a valid IKE phase-1 SA.<br>• Red indicates that IKE phase-1 SA is not available or has expired. |
| 4. View the **Tunnel Interface Status**.<br>• Green indicates that the tunnel interface is up.<br>• Red indicates that the tunnel interface is down, because tunnel monitoring is enabled and the status is down. |

To troubleshoot a VPN tunnel that is not yet up, see Interpret VPN Error Messages.

## Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel

You can enable, disable, refresh or restart an IKE gateway or VPN tunnel to make troubleshooting easier.

| Enable or Disable an IKE Gateway or Tunnel | |
|---|---|
| Enable or disable an IKE gateway. | 1. Select **Network > Network Profiles > IKE Gateways** and select the gateway you want to enable or disable.<br>2. At the bottom of the screen, click **Enable** or **Disable**. |
| Enable or disable an IPSec tunnel. | 1. Select **Network > IPSec Tunnels** and select the tunnel you want to enable or disable.<br>2. At the bottom of the screen, click **Enable** or **Disable**. |

The refresh and restart behaviors for an IKE gateway and IPSec tunnel are as follows:

| | Refresh | Restart |
|---|---|---|
| IKE Gateway (IKE Phase 1) | Updates the onscreen statistics for the selected IKE gateway.<br>Equivalent to issuing a second `show` command in the CLI (after an initial `show` command). | Restarts the selected IKE gateway.<br>**IKEv2**: Also restarts any associated child IPSec security associations (SAs).<br>**IKEv1**: Does not restart the associated IPSec SAs.<br>A restart is disruptive to all existing sessions.<br>Equivalent to issuing a `clear, test, show` command sequence in the CLI. |


© Palo Alto Networks, Inc.

| | Refresh | Restart |
|---|---|---|
| IPSec Tunnel (IKE Phase 2) | Updates the onscreen statistics for the selected IPSec tunnel.<br>Equivalent to issuing a second `show` command in the CLI (after an initial `show` command). | Restarts the IPSec tunnel.<br>A restart is disruptive to all existing sessions.<br>Equivalent to issuing a `clear, test, show` command sequence in the CLI. |

As the table above indicates, restarting an IKEv2 gateway has a result different from restarting an IKEv1 gateway.

| Refresh or Restart an IKE Gateway or IPSec Tunnel | |
|---|---|
| Refresh or restart an IKE gateway. | 1. Select **Network > IPSec Tunnels** and select the tunnel for the gateway you want to refresh or restart.<br>2. In the row for that tunnel, under the Status column, click **IKE Info**.<br>3. At the bottom of the IKE Info screen, click the action you want:<br>• **Refresh**—Updates the statistics on the screen.<br>• **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is recreated. |
| Refresh or restart an IPSec tunnel.<br>You might determine that the tunnel needs to be refreshed or restarted because you use the tunnel monitor to monitor the tunnel status, or you use an external network monitor to monitor network connectivity through the IPSec tunnel | 1. Select **Network > IPSec Tunnels** and select the tunnel you want to refresh or restart.<br>2. In the row for that tunnel, under the Status column, click **Tunnel Info**.<br>3. At the bottom of the Tunnel Info screen, click the action you want:<br>• **Refresh**—Updates the onscreen statistics.<br>• **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is recreated. |

# Test VPN Connectivity

## Test Connectivity

- Initiate IKE phase 1 by either pinging a host across the tunnel or using the following CLI command:

  **test vpn ike-sa gateway** *gateway_name*

- Then enter the following command to test if IKE phase 1 is set up:

  **show vpn ike-sa gateway** *gateway_name*
  In the output, check if the Security Association displays. If it does not, review the system log messages to interpret the reason for failure.

- Initiate IKE phase 2 by either pinging a host from across the tunnel or using the following CLI command:

  **test vpn ipsec-sa tunnel** *tunnel_name*

- Then enter the following command to test if IKE phase 1 is set up:

  **show vpn ipsec-sa tunnel** *tunnel_name*
  In the output, check if the Security Association displays. If it does not, review the system log messages to interpret the reason for failure.

- To view the VPN traffic flow information, use the following command:

  **show vpn-flow**

```
admin@PA-500> show vpn flow
total tunnels configured:             1
filter - type IPSec, state any

total IPSec tunnel configured:        1
total IPSec tunnel shown:             1

name               id      state    local-ip       peer-ip       tunnel-i/f
------------------------------------------------------------------------------
vpn-to-siteB       5       active    100.1.1.1     200.1.1.1      tunnel.41
```

 © Palo Alto Networks, Inc.

## Interpret VPN Error Messages

The following table lists some of the common VPN error messages that are logged in the system log.

**Table: Syslog Error Messages for VPN Issues**

| If error is this: | Try this: |
|---|---|
| `IKE phase-1 negotiation is failed as initiator, main mode. Failed SA: x.x.x.x[500]-y.y.y.y[500] cookie:84222f276c2fa2e9:0000000000000000 due to timeout.`<br>or<br>`IKE phase 1 negotiation is failed. Couldn't find configuration for IKE phase-1 request for peer IP x.x.x.x[1929]` | • Verify that the public IP address for each VPN peer is accurate in the IKE Gateway configuration.<br>• Verify that the IP addresses can be pinged and that routing issues are not causing the connection failure. |
| `Received unencrypted notify payload (no proposal chosen) from IP x.x.x.x[500] to y.y.y.y[500], ignored...`<br>or<br>`IKE phase-1 negotiation is failed. Unable to process peer's SA payload.` | Check the IKE Crypto profile configuration to verify that the proposals on both sides have a common encryption, authentication, and DH Group proposal. |
| `pfs group mismatched:my: 2peer: 0`<br>or<br>`IKE phase-2 negotiation failed when processing SA payload. No suitable proposal found in peer's SA payload.` | Check the IPSec Crypto profile configuration to verify that:<br>• pfs is either enabled or disabled on both VPN peers<br>• the DH Groups proposed by each peer has at least one DH Group in common |
| `IKE phase-2 negotiation failed when processing Proxy ID. Received local id x.x.x.x/x type IPv4 address protocol 0 port 0, received remote id y.y.y.y/y type IPv4 address protocol 0 port 0.` | The VPN peer on one end is using policy-based VPN. You must configure a Proxy ID on the Palo Alto Networks firewall. See Step 8. |

# Site-to-Site VPN Quick Configs

The following sections provide instructions for configuring some common VPN deployments:

- Site-to-Site VPN with Static Routing
- Site-to-Site VPN with OSPF
- Site-to-Site VPN with Static and Dynamic Routing


© Palo Alto Networks, Inc.

## Site-to-Site VPN with Static Routing

The following example shows a VPN connection between two sites that use static routes. Without dynamic routing, the tunnel interfaces on VPN Peer A and VPN Peer B do not require an IP address because the firewall automatically uses the tunnel interface as the next hop for routing traffic across the sites. However, to enable tunnel monitoring, a static IP address has been assigned to each tunnel interface.

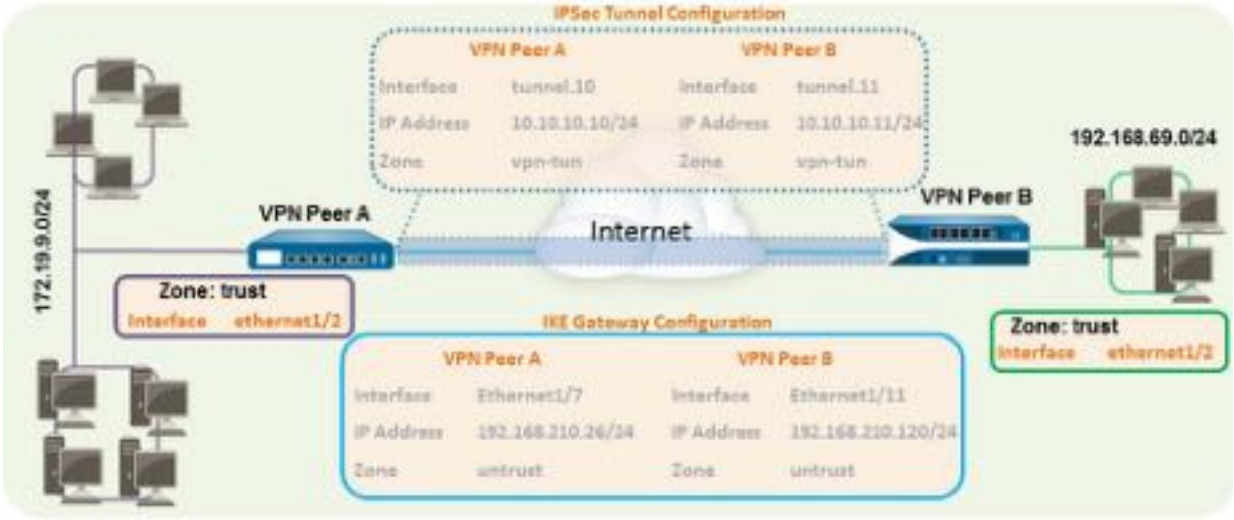

| Quick Config: Site-to-Site VPN with Static Routing | |
|---|---|
| Step 1 Configure a Layer 3 interface.<br><br>This interface is used for the IKE phase-1 tunnel. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—192.168.210.26/24<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—192.168.210.120/24 |

 © Palo Alto Networks, Inc.

**Quick Config: Site-to-Site VPN with Static Routing**

Step 2 Create a tunnel interface and attach it to a virtual router and security zone.

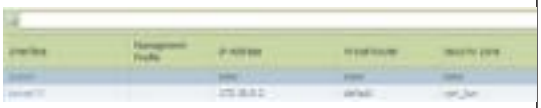

1. Select **Network > Interfaces > Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.1**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.
4. Select the **Virtual Router**.
5. (Optional) Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface.

   With static routes, the tunnel interface does not require an IP address. For traffic that is destined to a specified subnet/IP address, the tunnel interface will automatically become the next hop. Consider adding an IP address if you want to enable tunnel monitoring.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:
   - **Interface**—tunnel.11
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—172.19.9.2/24

   The configuration for VPN Peer B is:
   - **Interface**—tunnel.12
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—192.168.69.2/24

Step 3 Configure a static route, on the virtual router, to the destination subnet.

1. Select **Network > Virtual Router** and click the router you defined in the prior step.
2. Select **Static Route**, click **Add**, and enter a new route to access the subnet that is at the other end of the tunnel.

   In this example, the configuration for VPN Peer A is:
   - **Destination**—192.168.69.0/24
   - **Interface**—tunnel.11

   The configuration for VPN Peer B is:
   - **Destination**—172.19.9.0/24
   - **Interface**—tunnel.12

| Quick Config: Site-to-Site VPN with Static Routing | | |
|---|---|---|
| Step 4 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).<br><br>Complete this task on both peers and make sure to set identical values. | 1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.<br><br> <br><br>2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.<br><br> |
| Step 5 | Set up the IKE Gateway. | 1. Select **Network > Network Profiles > IKE Gateway.**<br>2. Click **Add** and configure the options in the **General** tab**.**<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Local IP address**—192.168.210.26/24<br>• **Peer IP type/address**—static/192.168.210.120<br>• **Preshared keys**—enter a value<br>• **Local identification**—None; this means that the local IP address will be used as the local identification value.<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Local IP address**—192.168.210.120/24<br>• **Peer IP type/address**—static/192.168.210.26<br>• **Preshared keys**—enter same value as on Peer A<br>• **Local identification**—None<br>3. Select **Advanced Phase 1 Options** and select the IKE Crypto profile you created earlier to use for IKE phase 1. |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static Routing | | |
|---|---|---|
| Step 6 | Set up the IPSec Tunnel.<br> 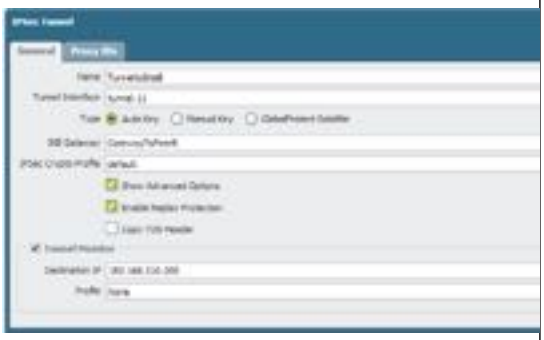 | 1. Select **Network > IPSec Tunnels**.<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Tunnel Interface**—tunnel.11<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IPSec Crypto profile defined in Step 4.<br>The configuration for VPN Peer B is:<br>• **Tunnel Interface**—tunnel.12<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IPSec Crypto defined in Step 4.<br>3. (Optional) Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity. Typically, the tunnel interface IP address for the VPN Peer is used.<br>4. (Optional) To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile. |
| Step 7 | Create policies to allow traffic between the sites (subnets). | 1. Select **Policies > Security**.<br>2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses. |
| Step 8 | Save any pending configuration changes. | Click **Commit**. |
| Step 9 | Test VPN connectivity. | See View the Status of the Tunnels. |

## Site-to-Site VPN with OSPF

In this example, each site uses OSPF for dynamic routing of traffic. The tunnel IP address on each VPN peer is statically assigned and serves as the next hop for routing traffic between the two sites.

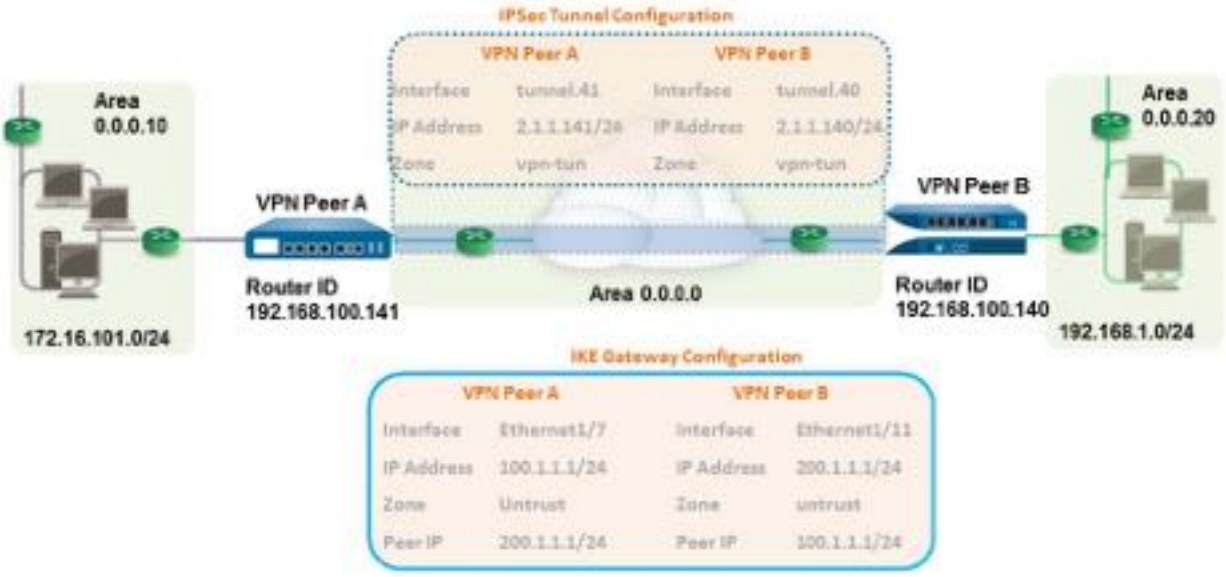


 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | |
|---|---|
| Step 1 Configure the Layer 3 interfaces on each firewall. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—100.1.1.1/24<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—200.1.1.1/24 |

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | |
|---|---|
| Step 2 Create a tunnel interface and attach it to a virtual router and security zone.<br> | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, say, **.11**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.<br>This IP address will be used as the next hop IP address to route traffic to the tunnel and can also be used to monitor the status of the tunnel.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—tunnel.41<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.141/24<br>The configuration for VPN Peer B is:<br>• **Interface**—tunnel.40<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.140/24 |
| Step 3 Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).<br>Complete this task on both peers and make sure to set identical values. | 1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.<br> <br>2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.<br> |


© Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | | |
|---|---|---|
| Step 4 | Set up the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall.<br><br>For more information on the OSPF options that are available on the firewall, see Configure OSPF.<br><br>Use Broadcast as the link type when there are more than two OSPF routers that need to exchange routing information. | 1. Select **Network > Virtual Routers**, and select the default router or add a new router.<br>2. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.<br>3. In this example, the OSPF configuration for VPN Peer A is:<br>– **Router ID**: 192.168.100.141<br>– **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p<br>– **Area ID**: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast<br>The OSPF configuration for VPN Peer B is:<br>– **Router ID**: 192.168.100.140<br>– **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p<br>– **Area ID**: 0.0.0.20 that is assigned to the interface Ethernet1/15 and Link Type: Broadcast |
| Step 5 | Set up the IKE Gateway.<br><br>This examples uses static IP addresses for both VPN peers. Typically, the corporate office uses a statically configured IP address, and the branch side can be a dynamic IP address; dynamic IP addresses are not best suited for configuring stable services such as VPN. | 1. Select **Network > Network Profiles > IKE Gateway.**<br>2. Click **Add** and configure the options in the **General** tab**.**<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Local IP address**—100.1.1.1/24<br>• **Peer IP address**—200.1.1.1/24<br>• **Preshared keys**—enter a value<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Local IP address**—200.1.1.1/24<br>• **Peer IP address**—100.1.1.1/24<br>• **Preshared keys**—enter same value as on Peer A<br>3. Select the IKE Crypto profile you created earlier to use for IKE phase 1. |

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | |
|---|---|
| Step 6 Set up the IPSec Tunnel. | 1. Select **Network > IPSec Tunnels**.<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Tunnel Interface**—tunnel.41<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>The configuration for VPN Peer B is:<br>• **Tunnel Interface**—tunnel.40<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.<br>4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile. |
| Step 7 Create policies to allow traffic between the sites (subnets). | 1. Select **Policies > Security**.<br>2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses. |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF | | |
|---|---|---|
| Step 8 | Verify OSPF adjacencies and routes from the CLI. | Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.<br>• **show routing protocol ospf neighbor**<br> 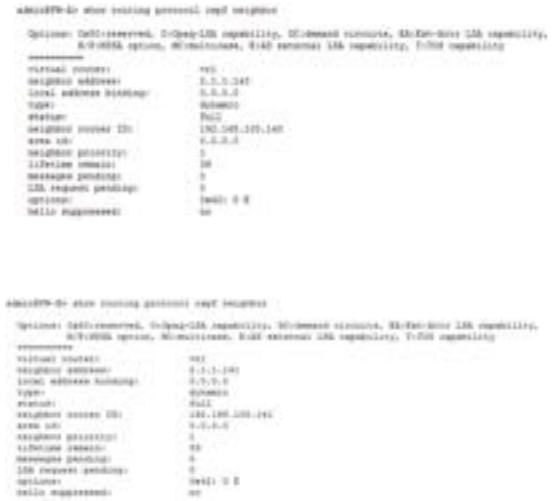 |
| | | • **show routing route type ospf**<br> 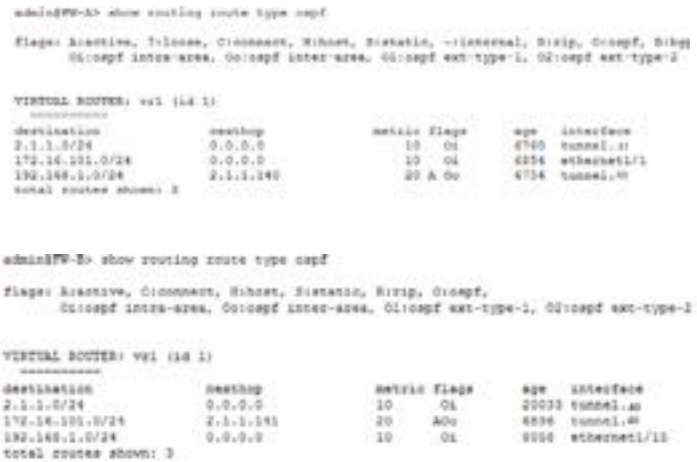 |
| Step 9 | Test VPN connectivity. | See Set Up Tunnel Monitoring and View the Status of the Tunnels. |

## Site-to-Site VPN with Static and Dynamic Routing

In this example, one site uses static routes and the other site uses OSPF. When the routing protocol is not the same between the locations, the tunnel interface on each firewall must be configured with a static IP address. Then, to allow the exchange of routing information, the firewall that participates in both the static and dynamic routing process must be configured with a *Redistribution profile*. Configuring the redistribution profile enables the virtual router to redistribute and filter routes between protocols—static routes, connected routes, and hosts—from the static autonomous system to the OSPF autonomous system. Without this redistribution profile, each protocol functions on its own and does not exchange any route information with other protocols running on the same virtual router.

In this example, the satellite office has static routes and all traffic destined to the 192.168.x.x network is routed to tunnel.41. The virtual router on VPN Peer B participates in both the static and the dynamic routing process and is configured with a redistribution profile in order to propagate (*export*) the static routes to the OSPF autonomous system.

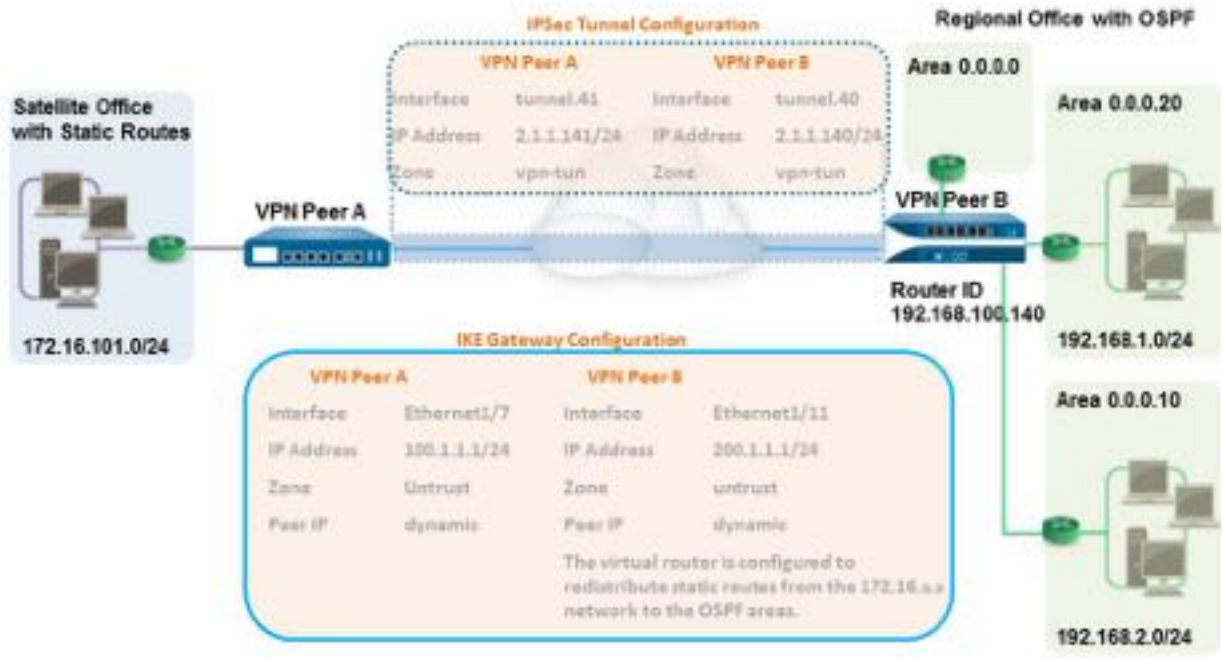


 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | | |
|---|---|---|
| Step 1 | Configure the Layer 3 interfaces on each firewall. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—100.1.1.1/24<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—untrust<br>• **Virtual Router**—default<br>• **IPv4**—200.1.1.1/24 |
| Step 2 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).<br>Complete this task on both peers and make sure to set identical values. | 1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.<br> <br>2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.<br> |

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 3 Set up the IKE Gateway.<br><br>With pre-shared keys, to add authentication scrutiny when setting up the IKE phase-1 tunnel, you can set up Local and Peer Identification attributes and a corresponding value that is matched in the IKE negotiation process. | 1. Select **Network > Network Profiles > IKE Gateway.**<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—ethernet1/7<br>• **Local IP address**—100.1.1.1/24<br>• **Peer IP type**—dynamic<br>• **Preshared keys**—enter a value<br>• **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer A.<br>• **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer B<br>The configuration for VPN Peer B is:<br>• **Interface**—ethernet1/11<br>• **Local IP address**—200.1.1.1/24<br>• **Peer IP address**—dynamic<br>• **Preshared keys**—enter same value as on Peer A<br>• **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer B<br>• **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer A<br>3. Select the IKE Crypto profile you created earlier to use for IKE phase 1. |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | | |
|---|---|---|
| Step 4 | Create a tunnel interface and attach it to a virtual router and security zone.<br> | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, say, **.41**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.<br>This IP address will be used to route traffic to the tunnel and to monitor the status of the tunnel.<br>6. To save the interface configuration, click **OK**.<br>In this example, the configuration for VPN Peer A is:<br>• **Interface**—tunnel.41<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.141/24<br>The configuration for VPN Peer B is:<br>• **Interface**—tunnel.42<br>• **Security Zone**—vpn_tun<br>• **Virtual Router**—default<br>• **IPv4**—2.1.1.140/24 |
| Step 5 | Specify the interface to route traffic to a destination on the 192.168.x.x network. | 1. On VPN Peer A, select the virtual router.<br>2. Select **Static Routes**, and **Add** tunnel.41 as the **Interface** for routing traffic with a **Destination** in the 192.168.x.x network. |

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 6 Set up the static route and the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall. | 1. On VPN Peer B, select **Network > Virtual Routers**, and select the default router or add a new router.<br>2. Select **Static Routes** and **Add** the tunnel IP address as the next hop for traffic in the 172.168.x.x. network.<br>Assign the desired route metric; using a lower the value makes the a higher priority for route selection in the forwarding table.<br>3. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.<br>4. In this example, the OSPF configuration for VPN Peer B is:<br>• Router ID: 192.168.100.140<br>• Area ID: 0.0.0.0 is assigned to the interface Ethernet 1/12 Link type: Broadcast<br>• Area ID: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast<br>• Area ID: 0.0.0.20 is assigned to the interface Ethernet1/15 and Link Type: Broadcast |
| Step 7 Create a redistribution profile to inject the static routes into the OSPF autonomous system.<br> <br> | 1. Create a redistribution profile on VPN Peer B.<br>a. Select **Network > Virtual Routers**, and select the router you used above.<br>b. Select **Redistribution Profiles,** and click **Add**.<br>c. Enter a Name for the profile and select **Redist** and assign a **Priority** value. If you have configured multiple profiles, the profile with the lowest priority value is matched first.<br>d. Set **Source Type** as **static**, and click **OK**. The static route defined in Step 6-2 will be used for the redistribution.<br>2. Inject the static routes in to the OSPF system.<br>a. Select **OSPF> Export Rules** (for IPv4) or **OSPFv3> Export Rules** (for IPv6).<br>b. Click **Add**, and select the redistribution profile that you just created.<br>c. Select how the external routes are brought into the OSPF system. The default option, **Ext2** calculates the total cost of the route using only the external metrics. To use both internal and external OSPF metrics, use **Ext1.**<br>d. Assign a **Metric** (cost value) for the routes injected into the OSPF system. This option allows you to change the metric for the injected route as it comes into the OSPF system.<br>e. Click **OK** to save the changes. |

 © Palo Alto Networks, Inc.

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | | |
|---|---|---|
| Step 8 | Set up the IPSec Tunnel. | 1. Select **Network > IPSec Tunnels**.<br>2. Click **Add** and configure the options in the **General** tab.<br>In this example, the configuration for VPN Peer A is:<br>• **Tunnel Interface**—tunnel.41<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>The configuration for VPN Peer B is:<br>• **Tunnel Interface**—tunnel.40<br>• **Type**—Auto Key<br>• **IKE Gateway**—Select the IKE Gateway defined above.<br>• **IPSec Crypto Profile**—Select the IKE Gateway defined above.<br>3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.<br>4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile. |
| Step 9 | Create policies to allow traffic between the sites (subnets). | 1. Select **Policies > Security**.<br>2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses. |

| Quick Config: Site-to-Site VPN with Static and Dynamic Routing | |
|---|---|
| Step 10 Verify OSPF adjacencies and routes from the CLI. | Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.<br>• **`show routing protocol ospf neighbor`**<br> 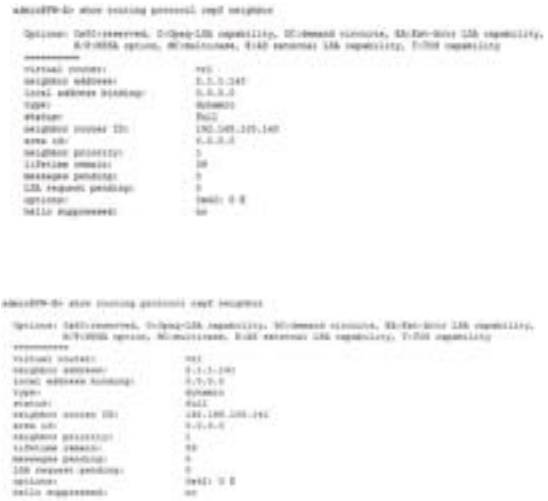 |
| | • **show routing route**<br>The following is an example of the output on each VPN peer.<br> |
| Step 11 Test VPN connectivity. | See Set Up Tunnel Monitoring and View the Status of the Tunnels. |


© Palo Alto Networks, Inc.

# Large Scale VPN (LSVPN)

The GlobalProtect Large Scale VPN (LSVPN) feature on the Palo Alto Networks next-generation firewall simplifies the deployment of traditional hub and spoke VPNs, enabling you to quickly deploy enterprise networks with several branch offices with a minimum amount of configuration required on the remote *satellite* devices. This solution uses certificates for device authentication and IPSec to secure data.

LSVPN enables site-to-site VPNs between Palo Alto Networks firewalls. To set up a site-to-site VPN between a Palo Alto Networks firewall and another device, see VPNs.

The following topics describe the LSVPN components and how to set them up to enable site-to-site VPN services between Palo Alto Networks firewalls:

- LSVPN Overview
- Create Interfaces and Zones for the LSVPN
- Enable SSL Between GlobalProtect LSVPN Components
- Configure the Portal to Authenticate Satellites
- Configure GlobalProtect Gateways for LSVPN
- Configure the GlobalProtect Portal for LSVPN
- Prepare the Satellite Device to Join the LSVPN
- Verify the LSVPN Configuration
- LSVPN Quick Configs

# LSVPN Overview

GlobalProtect provides a complete infrastructure for managing secure access to corporate resources from your remote sites. This infrastructure includes the following components:

- **GlobalProtect Portal**—Provides the management functions for your GlobalProtect LSVPN infrastructure. Every satellite that participates in the GlobalProtect LSVPN receives configuration information from the portal, including configuration information to enable the satellites (the spokes) to connect to the gateways (the hubs). You configure the portal on an interface on any Palo Alto Networks next-generation firewall.
- **GlobalProtect Gateways**—A Palo Alto Networks firewall that provides the tunnel end point for satellite connections. The resources that the satellites access is protected by security policy on the gateway. It is not required to have a separate portal and gateway; a single firewall can function both as portal and gateway.
- **GlobalProtect Satellite**—A Palo Alto Networks firewall at a remote site that establishes IPSec tunnels with the gateway(s) at your corporate office(s) for secure access to centralized resources. Configuration on the satellite firewall is minimal, enabling you to quickly and easily scale your VPN as you add new sites.

The following diagram illustrates how the GlobalProtect LSVPN components work together.

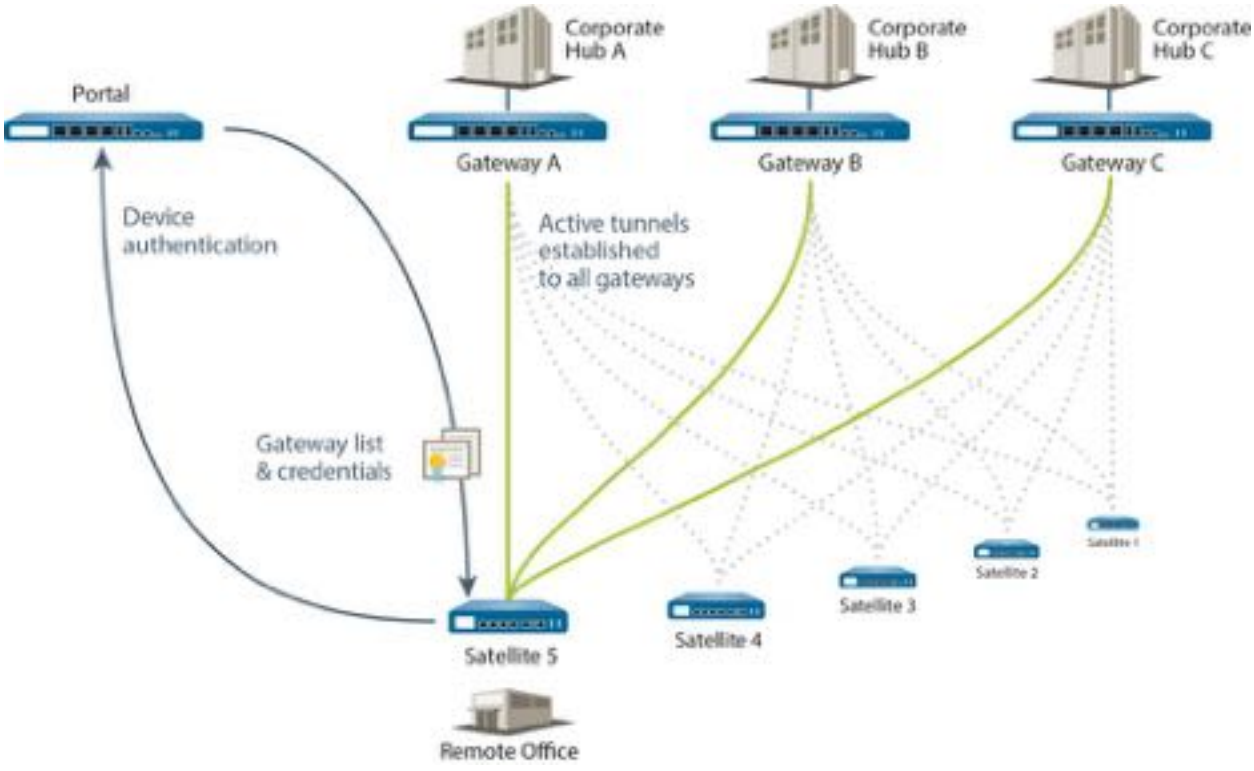


 © Palo Alto Networks, Inc.

# Create Interfaces and Zones for the LSVPN

You must configure the following interfaces and zones for your LSVPN infrastructure:

- **GlobalProtect portal**—Requires a Layer 3 interface for GlobalProtect satellites to connect to. If the portal and gateway are on the same firewall, they can use the same interface. The portal must be in a zone that is accessible from your branch offices.
- **GlobalProtect gateways**—Requires three interfaces: a Layer 3 interface in the zone that is reachable by the remote satellites, an internal interface in the trust zone that connects to the protected resources, and a logical tunnel interface for terminating the VPN tunnels from the satellites. Unlike other site-to-site VPN solutions, the GlobalProtect gateway only requires a single tunnel interface, which it will use for tunnel connections with all of your remote satellites (point-to-multi-point). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface.
- **GlobalProtect satellites**—Requires a single tunnel interface for establishing a VPN with the remote gateways (up to a maximum of 25 gateways). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface.

For more information about portals, gateways, and satellites see LSVPN Overview.

| Set Up Interfaces and Zones for the GlobalProtect LSVPN | |
|---|---|
| Step 1 Configure a Layer 3 interface.<br><br>The portal and each gateway and satellite all require a Layer 3 interface to enable traffic to be routed between sites.<br><br>If the gateway and portal are on the same firewall, you can use a single interface for both components.<br><br>IPv6 addresses are not supported with LSVPN. | 1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for GlobalProtect LSVPN.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, select the **Security Zone** to which the interface belongs:<br>• The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.<br>• If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.<br>4. Select the **Virtual Router** to use.<br>5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.11.100/24.<br>6. To save the interface configuration, click **OK**. |

| Set Up Interfaces and Zones for the GlobalProtect LSVPN (Continued) | | |
|---|---|---|
| Step 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.<br><br>IP addresses are not required on the tunnel interface unless plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.<br><br>Make sure to enable User-ID in the zone where the VPN tunnels terminate. | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:<br>• To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.<br>• (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *lsvpn-tun*), select the **Enable User Identification** check box, and then click **OK**.<br>4. Select the **Virtual Router**.<br>5. (Optional) If you want to assign an IP address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.11.33/24.<br>6. To save the interface configuration, click **OK**. |
| Step 3 | If you created a separate zone for tunnel termination of VPN connections, create a security policy to enable traffic flow between the VPN zone and your trust zone. | For example, the following policy rule enables traffic between the *lsvpn-tun* zone and the *L3-Trust* zone.<br> 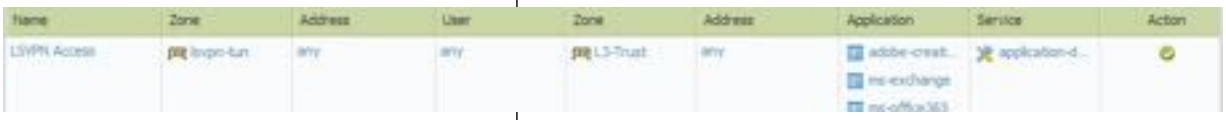 |
| Step 4 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# Enable SSL Between GlobalProtect LSVPN Components

All interaction between the GlobalProtect components occurs over an SSL/TLS connection. Therefore, you must generate and/or install the required certificates before configuring each component so that you can reference the appropriate certificate(s) and/or certificate profiles in the configurations for each component. The following sections describe the supported methods of certificate deployment, descriptions and best practice guidelines for the various GlobalProtect certificates, and provide instructions for generating and deploying the required certificates:


- About Certificate Deployment
- Deploy Server Certificates to the GlobalProtect LSVPN Components


## About Certificate Deployment

There are two basic approaches to deploying certificates for GlobalProtect LSVPN:

- **Enterprise Certificate Authority**—If you already have your own enterprise certificate authority, you can use this internal CA to issue an intermediate CA certificate for the GlobalProtect portal to enable it to issue certificates to the GlobalProtect gateways and satellites.
- **Self-Signed Certificates**—You can generate a self-signed root CA certificate on the firewall and use it to issue server certificates for the portal, gateway(s), and satellite(s). As a best practice, create a self-signed root CA certificate on the portal and use it to issue server certificates for the gateways and satellites. This way, the private key used for certificate signing stays on the portal.

## Deploy Server Certificates to the GlobalProtect LSVPN Components

The GlobalProtect LSVPN components use SSL/TLS to mutually authenticate. Before deploying the LSVPN, you must assign an SSL/TLS service profile to each portal and gateway. The profile specifies the server certificate and allowed TLS versions for communication with satellite devices. You don't need to create SSL/TLS service profiles for the satellite devices because the portal will issue a server certificate for each device during the first connection as part of the device registration process.

In addition, you must import the root certificate authority (CA) certificate used to issue the server certificates onto each firewall that you plan to host as a gateway or satellite device. Finally, on each gateway and satellite device participating in the LSVPN, you must configure a certificate profile that will enable them to establish an SSL/TLS connection using mutual authentication.

The following workflow shows the best practice steps for deploying SSL certificates to the GlobalProtect LSVPN components:

© Palo Alto Networks, Inc.

| Deploy SSL Server Certificates to the GlobalProtect Components | | |
|---|---|---|
| Step 1 | On the firewall hosting the GlobalProtect portal, create the root CA certificate for signing the certificates of the GlobalProtect components. | Create a Self-Signed Root CA Certificate:<br>1. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Generate**.<br>2. Enter a **Certificate Name**, such as `LSVPN_CA`.<br>3. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed).<br>4. Select the **Certificate Authority** check box and then click **OK** to generate the certificate. |
| Step 2 | Create SSL/TLS service profiles for the GlobalProtect portal and gateways.<br><br>For the portal and each gateway, you must assign an SSL/TLS service profile that references a unique self-signed server certificate.<br><br> The best practice is to issue all of the required certificates on the portal, so that the signing certificate (with the private key) doesn't have to be exported.<br><br> If the GlobalProtect portal and gateway are on the same firewall interface, you can use the same server certificate for both components. | 1. Use the root CA on the portal to Generate a Certificate on the Device for each gateway you will deploy:<br>a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Generate**.<br>b. Enter a **Certificate Name**.<br>c. Enter the FQDN (recommended) or IP address of the interface where you plan to configure the gateway in the **Common Name** field.<br>d. In the **Signed By** field, select the `LSVPN_CA` certificate you just created.<br>e. In the Certificate Attributes section, click **Add** and define the attributes to uniquely identify the gateway. If you add a **Host Name** attribute (which populates the SAN field of the certificate), it must exactly match the value you defined for the **Common Name**.<br>f. **Generate** the certificate.<br>2. Configure an SSL/TLS Service Profile for the portal and each gateway:<br>a. Select **Device > Certificate Management > SSL/TLS Service Profile** and click **Add**.<br>b. Enter a **Name** to identify the profile and select the server **Certificate** you just created for the portal or gateway.<br>c. Define the range of TLS versions (**Min Version** to **Max Version**) allowed for communicating with satellite devices and click **OK**. |

 © Palo Alto Networks, Inc.

| Deploy SSL Server Certificates to the GlobalProtect Components (Continued) | |
|---|---|
| Step 3 Deploy the self-signed server certificates to the gateways.<br><br>**Best Practices:**<br>•Export the self-signed server certificates issued by the root CA from the portal and import them onto the gateways.<br>•Be sure to issue a unique server certificate for each gateway.<br>•The Common Name (CN) and, if applicable, the Subject Alternative Name (SAN) fields of the certificate must match the IP address or fully qualified domain name (FQDN) of the interface where you configure the gateway. | 1. On the portal, select **Device > Certificate Management > Certificates > Device Certificates**, select the gateway certificate you want to deploy, and click **Export**.<br>2. Select **Encrypted Private Key and Certificate (PKCS12)** from the **File Format** drop-down.<br>3. Enter (and re-enter) a **Passphrase** to encrypt the private key associated with the certificate and then click **OK** to download the PKCS12 file to your computer.<br>4. On the gateway, select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>5. Enter a **Certificate Name**.<br>6. Enter the path and name to the **Certificate File** you just downloaded from the portal, or **Browse** to find the file.<br>7. Select **Encrypted Private Key and Certificate (PKCS12)** as the **File Format**.<br>8. Enter the path and name to the PKCS12 file in the **Key File** field or **Browse** to find it.<br>9. Enter and re-enter the **Passphrase** you used to encrypt the private key when you exported it from the portal and then click **OK** to import the certificate and key. |
| Step 4 Import the root CA certificate used to issue server certificates for the LSVPN components.<br><br>You must import the root CA certificate onto all gateways and satellites. For security reasons, make sure you export the certificate only, and not the associated private key. | 1. Download the root CA certificate from the portal.<br>a. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>b. Select the root CA certificate used to issue certificates for the LSVPN components and click **Export**.<br>c. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (Do not export the private key.)<br>2. On the firewalls hosting the gateways and satellites, import the root CA certificate.<br>a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.<br>c. **Browse** to the **Certificate File** you downloaded from the CA.<br>d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.<br>e. Select the certificate you just imported on the **Device Certificates** tab to open it.<br>f. Select **Trusted Root CA** and then click **OK**.<br>g. **Commit** the changes. |

| Deploy SSL Server Certificates to the GlobalProtect Components (Continued) | | |
|---|---|---|
| Step 5 | Create a certificate profile.<br>The GlobalProtect LSVPN portal and each gateway require a certificate profile that specifies which certificate to use to authenticate the satellite devices. | 1. Select **Device > Certificate Management > Certificate Profile** and click **Add** and enter a profile **Name**.<br>2. Make sure **Username Field** is set to **None**.<br>3. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you imported in Step 4.<br>4. (Optional, but recommended) Enable use of CRL and/or OCSP to enable certificate status verification.<br>5. Click **OK** to save the profile. |
| Step 6 | Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

# Configure the Portal to Authenticate Satellites

In order to register with the LSVPN, each satellite must establish an SSL/TLS connection with the portal. After establishing the connection, the portal authenticates the satellite device to ensure that is authorized to join the LSVPN. After successfully authenticating the satellite, the portal will issue a server certificate for the satellite and push the LSVPN configuration specifying the gateways to which the satellite can connect and the root CA certificate required to establish an SSL connection with the gateways.

There are two ways that the satellite can authenticate to the portal during its initial connection:

- **Serial number**—You can configure the portal with the serial number of the satellite firewalls that are authorized to join the LSVPN. During the initial satellite connection to the portal, the satellite presents its serial number to the portal and if the portal has the serial number in its configuration, the satellite will be successfully authenticated. You add the serial numbers of authorized satellites when you configure the portal. See Configure the Portal.
- **Username and password**—If you would rather provision your satellites without manually entering the serial numbers of the satellite devices into the portal configuration, you can instead require the satellite administrator to authenticate when establishing the initial connection to the portal. Although the portal will always look for the serial number in the initial request from the satellite, if it cannot identify the serial number, the satellite administrator must provide a username and password to authenticate to the portal. Because the portal will always fall back to this form of authentication, you must create an authentication profile in order to commit the portal configuration. This requires that you set up an authentication profile for the portal LSVPN configuration even if you plan to authenticate satellites using the serial number.

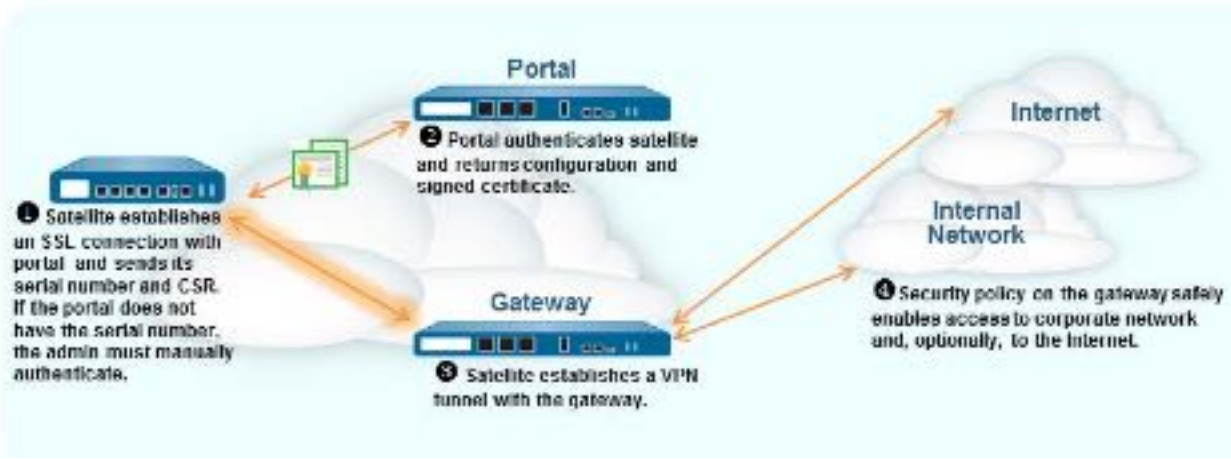


The following workflow describes how to set up the portal to authenticate satellites against an existing authentication service. GlobalProtect LSVPN supports external authentication using a local database, LDAP (including Active Directory), Kerberos, TACACS+, or RADIUS.

| Set Up Satellite Authentication | | |
|---|---|---|
| Step 1 | (External authentication only) Create a server profile on the portal.<br><br>The server profile defines how the firewall connects to an external authentication service to validate the authentication credentials that the satellite device administrator enters.<br><br> If you use local authentication, skip this step and instead add a local user for the satellite device administrator: see Create the local database. | Configure a server profile for the authentication service type:<br>• Configure a RADIUS Server Profile.<br>• Configure a TACACS+ Server Profile.<br>• Configure an LDAP Server Profile. If you use LDAP to connect to Active Directory (AD), create a separate LDAP server profile for every AD domain.<br>• Configure a Kerberos Server Profile. |
| Step 2 | Configure an authentication profile.<br><br>The authentication profile defines which server profile to use to authenticate satellite devices. | 1. Select **Device > Authentication Profile** and click **Add**.<br>2. Enter a **Name** for the profile and then select the authentication **Type**. If the **Type** is an external service, select the **Server Profile** you created in Step 1. If you added a local user instead, set the **Type** to **Local Database**.<br>3. Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# Configure GlobalProtect Gateways for LSVPN

Because the GlobalProtect configuration that the portal delivers to the satellites includes the list of gateways the satellite can connect to, it is a good idea to configure the gateways before configuring the portal.


- ▲ Prerequisite Tasks
- ▲ Configure the Gateway


## Prerequisite Tasks

Before you can configure the GlobalProtect gateway, you must complete the following tasks:

- ❒ Create Interfaces and Zones for the LSVPN on the interface where you will configure each gateway. You must configure both the physical interface and the virtual tunnel interface.
- ❒ Enable SSL Between GlobalProtect LSVPN Components by configuring the gateway server certificates, SSL/TLS service profiles, and certificate profile required to establish a mutual SSL/TLS connection from the GlobalProtect satellite devices to the gateway.

## Configure the Gateway

After you have completed the Prerequisite Tasks, configure each GlobalProtect gateway to participate in the LSVPN as follows:

| Configure the Gateway for LSVPN | | |
|---|---|---|
| Step 1 | Add a gateway. | 1. Select **Network > GlobalProtect > Gateways** and click **Add**.<br>2. On the **General** tab, enter a **Name** for the gateway. The gateway name should not contain any spaces and as a best practice it should include the location or other descriptive information that will help identify the gateway.<br>3. (Optional) Select the virtual system to which this gateway belongs from the **Location** field. |
| Step 2 | Specify the network information to enable satellite devices to connect to the gateway.<br><br>If you haven't yet created the network interface for the gateway, see Create Interfaces and Zones for the LSVPN for instructions. If you haven't yet created an SSL/TLS Service profile for the gateway, see Deploy Server Certificates to the GlobalProtect LSVPN Components. | 1. Select the **Interface** that satellite devices will use for ingress access to the gateway.<br>2. Select the **IP Address** for gateway access.<br>3. Select the **SSL/TLS Service Profile** for the gateway. |

© Palo Alto Networks, Inc.

| Configure the Gateway for LSVPN (Continued) | | |
|---|---|---|
| Step 3 | Select the certificate profile for the gateway to use to authenticate satellite devices attempting to establish tunnels.<br><br>If you have not yet set up the certificate profile, see Enable SSL Between GlobalProtect LSVPN Components for instructions. | Select the **Certificate Profile** to you created for SSL communication between the LSVPN components. |
| Step 4 | Configure the tunnel parameters and enable tunneling. | 1. On the GlobalProtect Gateway dialog, select **Satellite Configuration > Tunnel Settings**.<br>2. Select the **Tunnel Configuration** check box to enable tunneling.<br>3. Select the **Tunnel Interface** you defined in Step 2 in Create Interfaces and Zones for the LSVPN.<br>4. (Optional) If you want to preserve the Type of Service (ToS) information in the encapsulated packets, select the **Copy TOS** check box. |
| Step 5 | (Optional) Enable tunnel monitoring.<br><br>Tunnel monitoring enables satellite devices to monitor its gateway tunnel connection, allowing it to failover to a backup gateway if the connection fails. Failover to another gateway is the only type of tunnel monitoring profile supported with LSVPN. | 1. Select the **Tunnel Monitoring** check box.<br>2. Specify the **Destination IP** address the satellite devices should use to determine if the gateway is active. Alternatively, if you configured an IP address for the tunnel interface, you can leave this field blank and the tunnel monitor will instead use the tunnel interface to determine if the connection is active.<br>3. Select **Failover** from the **Tunnel Monitor Profile** drop-down (this is the only supported tunnel monitor profile for LSVPN). |
| Step 6 | Select the IPSec Crypto profile to use when establishing tunnel connections.<br><br>The profile specifies the type of IPSec encryption and the authentication method for securing the data that will traverse the tunnel. Because both tunnel endpoints in an LSVPN are trusted firewalls within your organization, you can typically use the default (predefined) profile, which uses ESP as the IPSec protocol, group2 for the DH group, AES-128-CBC for encryption, and SHA-1 for authentication. | In the **IPSec Crypto Profile** drop-down, select **default** to use the predefined profile or select **New IPSec Crypto Profile** to define a new profile. For details on the authentication and encryption options, see Define IPSec Crypto Profiles. |

 © Palo Alto Networks, Inc.

| Configure the Gateway for LSVPN (Continued) | | |
|---|---|---|
| Step 7 | Configure the network settings to assign the satellite devices during establishment of the IPSec tunnel.<br><br>You can also configure the satellite device to push the DNS settings to its local clients by configuring a DHCP server on the firewall hosting the satellite. In this configuration, the satellite will push DNS settings it learns from the gateway to the DHCP clients. | 1. On the GlobalProtect Gateway dialog, select **Satellite Configuration > Network Settings**.<br>2. (Optional) If clients local to the satellite device need to resolve FQDNs on the corporate network, configure the gateway to push DNS settings to the satellites in one of the following ways:<br>• Manually define the **Primary DNS**, **Secondary DNS**, and **DNS Suffix** settings to push to the satellites.<br>• If the gateway has an interface that is configured as a DHCP client, you can set the **Inheritance Source** to that interface and the GlobalProtect satellites will be assigned the same settings received by the DHCP client.<br>3. To specify the **IP Pool** of addresses to assign the tunnel interface on the satellite devices when the VPN is established, click **Add** and then specify the IP address range(s) to use.<br>4. To define what destination subnets to route through the tunnel click **Add** in the **Access Route** area and then enter the routes as follows:<br>• If you want to route all traffic from the satellites through the tunnel, leave this field blank. Note that in this case, all traffic except traffic destined for the local subnet will be tunneled to the gateway.<br>• To route only some traffic through the gateway (called *split tunneling*), specify the destination subnets that must be tunneled. In this case, the satellite will route traffic that is not destined for a specified access route using its own routing table. For example, you may choose to only tunnel traffic destined for your corporate network, and use the local satellite to safely enable Internet access.<br>• If you want to enable routing between satellites, enter the summary route for the network protected by each satellite. |
| Step 8 | (Optional) Define what routes, if any, the gateway will accept from satellites.<br><br>By default, the gateway will not add any routes satellites advertise to its routing table. If you do not want the gateway to accept routes from satellites, you do not need to complete this step. | 1. To enable the gateway to accept routes advertised by satellites, select **Satellite Configuration > Route Filter**.<br>2. Select the **Accept published routes** check box.<br>3. To filter which of the routes advertised by the satellites to add to the gateway routing table, click **Add** and then define the subnets to include. For example, if all the satellites are configured with subnet 192.168.x.0/24 on the LAN side, configuring a permitted route of 192.168.0.0/16 to enable the gateway to only accept routes from the satellite if it is in the 192.168.0.0/16 subnet. |
| Step 9 | Save the gateway configuration. | 1. Click **OK** to save the settings and close the GlobalProtect Gateway dialog.<br>2. **Commit** the configuration. |

# Configure the GlobalProtect Portal for LSVPN

The GlobalProtect portal provides the management functions for your GlobalProtect LSVPN. Every satellite system that participates in the LSVPN receives configuration information from the portal, including information about available gateways as well as the certificate it needs in order to connect to the gateways.

The following sections provide procedures for setting up the portal:

- Prerequisite Tasks
- Configure the Portal
- Define the Satellite Configurations

## Prerequisite Tasks

Before configuring the GlobalProtect portal, you must complete the following tasks:

- Create Interfaces and Zones for the LSVPN on the interface where you will configure the portal.
- Enable SSL Between GlobalProtect LSVPN Components by creating an SSL/TLS service profile for the portal server certificate, issuing gateway server certificates, and configuring the portal to issue server certificates for the GlobalProtect satellite devices.
- Configure the Portal to Authenticate Satellites by defining the authentication profile that the portal will use to authenticate satellite devices if the serial number is not available.
- Configure GlobalProtect Gateways for LSVPN.

## Configure the Portal

After you have completed the Prerequisite Tasks, configure the GlobalProtect portal as follows:

| **Configure the Portal for LSVPN** | |
|---|---|
| Step 1 Add the portal. | 1. Select **Network > GlobalProtect > Portals** and click **Add**.<br>2. On the **Portal Configuration** tab, enter a **Name** for the portal. The portal name should not contain any spaces.<br>3. (Optional) Select the virtual system to which this portal belongs from the **Location** field. |


© Palo Alto Networks, Inc.

| Configure the Portal for LSVPN (Continued) | | |
|---|---|---|
| Step 2 | Specify the network information to enable satellite devices to connect to the portal.<br><br>If you haven't yet created the network interface for the portal, see Create Interfaces and Zones for the LSVPN for instructions. If you haven't yet created an SSL/TLS service profile for the portal and issued gateway certificates, see Deploy Server Certificates to the GlobalProtect LSVPN Components. | 1. Select the **Interface** that satellite devices will use for ingress access to the portal.<br>2. Select the **IP Address** for satellite device access to the portal.<br>3. Select the **SSL/TLS Service Profile** to enable the satellite device to establish an SSL/TLS connection to the portal. |
| Step 3 | Specify an authentication profile for authenticating satellite devices.<br><br>If the portal can't validate the serial numbers of connecting satellite devices, it will fall back to the authentication profile. Therefore, you must configure an authentication profile before you can save the portal configuration (by clicking **OK**). | • Select the **Authentication Profile** you defined for authenticating satellite devices.<br>• If you haven't yet configured the authentication profile, select **New Authentication Profile** to Configure an authentication profile. |
| Step 4 | Continue with defining the configurations to push to the satellite devices or, if you have already created the satellite device configurations, save the portal configuration. | Click **OK** to save the portal configuration or continue to Define the Satellite Configurations. |

## Define the Satellite Configurations

When a GlobalProtect satellite connects and successfully authenticates to the GlobalProtect portal, the portal delivers a satellite configuration, which specifies what gateways the satellite can connect to. If all your satellites will use the same gateway and certificate configurations, you can create a single satellite configuration to deliver to all satellites upon successful authentication. However, if you require different satellite configurations—for example if you want one group of satellites to connect to one gateway and another group of satellites to connect to a different gateway—you can create a separate satellite configuration for each. The portal will then use the enrollment username/group name or the serial number of the satellite device to determine which satellite configuration to deploy. As with security rule evaluation, the portal looks for a match starting from the top of the list. When it finds a match, it delivers the corresponding configuration to the satellite.

For example, the following figure shows a network in which some branch offices require VPN access to the corporate applications protected by your perimeter firewalls and another site needs VPN access to the datacenter.

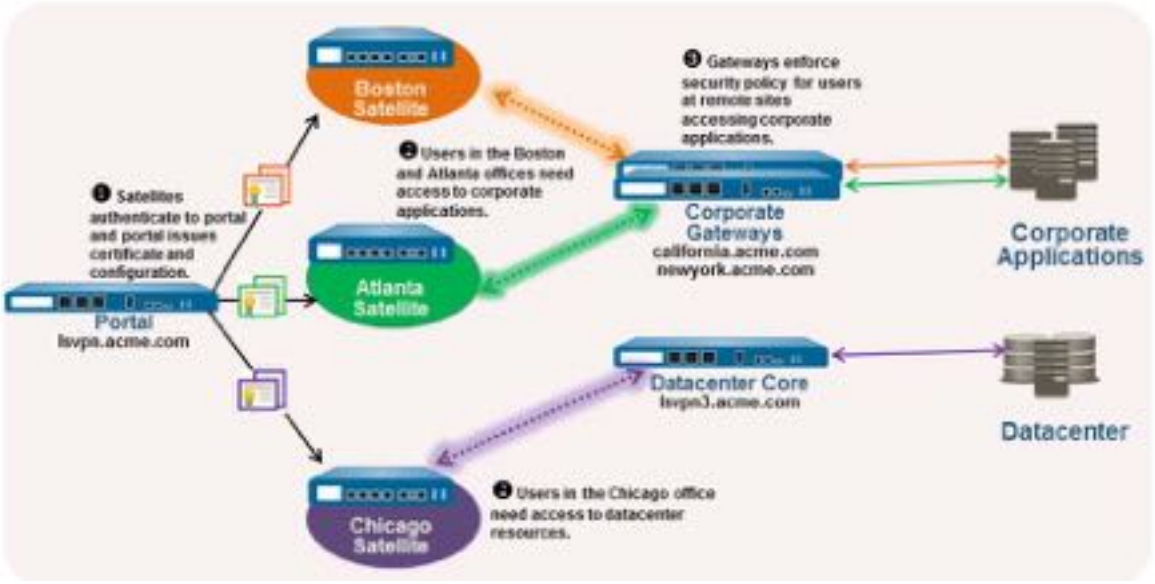


Use the following procedure to create one or more satellite configurations.

| Create a GlobalProtect Satellite Configuration | |
|---|---|
| Step 1 Specify the certificates required to enable satellites to participate in the LSVPN. | 1. Select **Network > GlobalProtect > Portals** and select the portal configuration for which you want to add a satellite configuration and then select the **Satellite Configuration** tab.<br>2. In the **Trusted Root CA** field, click **Add** and then select the CA certificate used to issue the gateway server certificates. The portal will deploy the root CA certificate you add here to all satellites as part of the configuration to enable the satellite to establish an SSL connection with the gateways. As a best practice, all of your gateways should use the same issuer.<br>If the root CA certificate used to issue your gateway server certificates is not on the portal, you can **Import** it now. See Enable SSL Between GlobalProtect LSVPN Components for details on how to import a root CA certificate.<br>3. Select the Root CA certificate that the portal will use to issue certificates to satellites upon successfully authenticating them from the **Issuing Certificate** drop-down. |
| Step 2 Add a satellite configuration.<br>The satellite configuration specifies the GlobalProtect LSVPN configuration settings to deploy to the connecting satellites. You must define at least one satellite configuration. | In the Satellite Configuration section, click **Add** and enter a **Name** for the configuration.<br>If you plan to create multiple configurations, make sure the name you define for each is descriptive enough to allow you to distinguish them. |

 © Palo Alto Networks, Inc.

| Create a GlobalProtect Satellite Configuration (Continued) | | |
|---|---|---|
| Step 3 | Specify which satellites to deploy this configuration to.<br><br>The portal uses the **Enrollment User/User Group** settings and/or **Devices** serial numbers to match a satellite to a configuration. Therefore, if you have multiple configurations, be sure to order them properly. As soon as the portal finds a match, it will deliver the configuration. Therefore, more specific configurations must precede more general ones. See Step 6 for instructions on ordering the list of satellite configurations. | Specify the match criteria for the satellite configuration as follows:<br>• To restrict this configuration to satellite devices with specific serial numbers, select the **Devices** tab, click **Add**, and enter serial number (you do not need to enter the satellite hostname; it will be automatically added when the satellite connects). Repeat this step for each satellite you want to receive this configuration.<br>• Select the **Enrollment User/User Group** tab, click **Add**, and then select the user or group you want to receive this configuration. Satellites that do not match on serial number will be required to authenticate as a user specified here (either an individual user or group member).<br><br>Before you can restrict the configuration to specific groups, you must Map Users to Groups. |
| Step 4 | Specify the gateways that satellites with this configuration can establish VPN tunnels with.<br><br>Routes published by the gateway are installed on the satellite as static routes. The metric for the static route is 10x the routing priority. If you have more than one gateway, make sure to also set the routing priority to ensure that routes advertised by backup gateways have higher metrics compared to the same routes advertised by primary gateways. For example, if you set the routing priority for the primary gateway and backup gateway to 1 and 10 respectively, the satellite will use 10 as the metric for the primary gateway and 100 as the metric for the backup gateway. | 1. On the **Gateways** tab, click **Add**.<br>2. Enter a descriptive **Name** for the gateway. The name you enter here should match the name you defined when you configured the gateway and should be descriptive enough identify the location of the gateway.<br>3. Enter the FQDN or IP address of the interface where the gateway is configured in the **Gateways** field. The address you specify must exactly match the Common Name (CN) in the gateway server certificate.<br>4. (Optional) If you are adding two or more gateways to the configuration, the **Routing Priority** helps the satellite pick the preferred gateway. Enter a value in the range of 1-25, with lower numbers having the higher priority (that is, the gateway the satellite will connect to if all gateways are available). The satellite will multiply the routing priority by 10 to determine the routing metric. |
| Step 5 | Save the satellite configuration. | 1. Click **OK** to save the satellite configuration.<br>2. If you want to add another satellite configuration, repeat Step 2 through Step 5. |
| Step 6 | Arrange the satellite configurations so that the proper configuration is deployed to each satellite. | • To move a satellite configuration up on the list of configurations, select the configuration and click **Move Up**.<br>• To move a satellite configuration down on the list of configurations, select the configuration and click **Move Down**. |
| Step 7 | Save the portal configuration. | 1. Click **OK** to save the settings and close the GlobalProtect Portal dialog.<br>2. **Commit** your changes. |

# Prepare the Satellite Device to Join the LSVPN

In order to participate in the LSVPN, the satellite devices require a minimal amount of configuration. Because the required configuration is minimal, you can pre-configure the devices before shipping them to your branch offices for installation.

| Prepare the Satellite Device to Join the GlobalProtect LSVPN | | |
|---|---|---|
| Step 1 | Configure a Layer 3 interface. | This is the physical interface the satellite will use to connect to the portal and the gateway. This interface must be in a zone that allows access outside of the local trust network. As a best practice, create a dedicated zone for VPN connections for visibility and control over traffic destined for the corporate gateways. |
| Step 2 | Configure the logical tunnel interface for the tunnel to use to establish VPN tunnels with the GlobalProtect gateways.<br> IP addresses are not required on the tunnel interface unless you plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues. | 1. Select **Network > Interfaces > Tunnel** and click **Add**.<br>2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select an existing zone or create a separate zone for VPN tunnel traffic by clicking **New Zone** and defining a **Name** for new zone (for example *lsvpnsat*).<br>4. In the **Virtual Router** drop-down, select **default**.<br>5. (Optional) If you want to assign an IP address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 2.2.2.11/24.<br>6. To save the interface configuration, click **OK**. |

 © Palo Alto Networks, Inc.

| Prepare the Satellite Device to Join the GlobalProtect LSVPN (Continued) | | |
|---|---|---|
| Step 3 | If you generated the portal server certificate using a Root CA that is not trusted by the satellite devices (for example, if you used self-signed certificates), import the root CA certificate used to issue the portal server certificate.<br><br>The root CA certificate is required to enable the satellite device to establish the initial connection with the portal to obtain the LSVPN configuration. | 1. Download the CA certificate that was used to generate the portal server certificates. If you are using self-signed certificates, export the root CA certificate from the portal as follows:<br>a. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>b. Select the CA certificate, and click **Export**.<br>c. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (You do not need to export the private key.)<br>2. Import the root CA certificate you just exported onto each satellite device as follows.<br>a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.<br>b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.<br>c. **Browse** to the **Certificate File** you downloaded from the CA.<br>d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.<br>e. Select the certificate you just imported on the **Device Certificates** tab to open it.<br>f. Select **Trusted Root CA** and then click **OK**. |
| Step 4 | Configure the IPSec tunnel configuration. | 1. Select **Network > IPSec Tunnels** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the IPSec configuration.<br>3. Select the **Tunnel Interface** you created for the satellite device.<br>4. Select **GlobalProtect Satellite** as the **Type**.<br>5. Enter the IP address or FQDN of the portal as the **Portal Address**.<br>6. Select the Layer 3 **Interface** you configured for the satellite device.<br>7. Select the **Local IP Address** to use on the selected interface. |

| Prepare the Satellite Device to Join the GlobalProtect LSVPN (Continued) | |
|---|---|
| Step 5 (Optional) Configure the satellite to publish local routes to the gateway.<br><br>Pushing routes to the gateway enables traffic to the subnets local to the satellite via the gateway. However, you must also configure the gateway to accept the routes as detailed in Step 8 in Configure the Gateway. | 1. To enable the satellite to push routes to the gateway, on the **Advanced** tab select **Publish all static and connected routes to Gateway**.<br><br>If you select this check box, the firewall will forward all static and connected routes from the satellite to the gateway. However, to prevent the creation of routing loops, the firewall will apply some route filters, such as the following:<br>• Default routes<br>• Routes within a virtual router other than the virtual router associated with the tunnel interface<br>• Routes using the tunnel interface<br>• Routes using the physical interface associated with the tunnel interface<br>2. (Optional) If you only want to push routes for specific subnets rather than all routes, click **Add** in the Subnet section and specify which subnet routes to publish. |
| Step 6 Save the satellite configuration. | 1. Click **OK** to save the IPSec tunnel settings.<br>2. Click **Commit**. |
| Step 7 If required, provide the credentials to allow the satellite to authenticate to the portal.<br><br>This step is only required if the portal was unable to find a serial number match in its configuration or if the serial number didn't work. In this case, the satellite will not be able to establish the tunnel with the gateway(s). | 1. Select **Network > IPSec Tunnels** and click the **Gateway Info** link in the Status column of the tunnel configuration you created for the LSVPN.<br>2. Click the **enter credentials** link in the **Portal Status** field and username and password required to authenticate the satellite to the portal.<br><br>After the portal successfully authenticates to the portal, it will receive its signed certificate and configuration, which it will use to connect to the gateway(s). You should see the tunnel establish and the **Status** change to **Active**.<br><br> |


© Palo Alto Networks, Inc.

# Verify the LSVPN Configuration

After configuring the portal, gateways, and satellite devices, verify that the satellites are able to connect to the portal and gateway and establish VPN tunnels with the gateway(s).

| Verify the LSVPN Configuration | | |
|---|---|---|
| Step 1 | Verify satellite connectivity with portal. | From the firewall hosting the portal, verify that satellites are successfully connecting by selecting **Network > GlobalProtect > Portal** and clicking **Satellite Info** in the Info column of the portal configuration entry.<br> |
| Step 1 | Verify satellite connectivity with the gateway(s). | On each firewall hosting a gateway, verify that satellites are able to establish VPN tunnels by selecting **Network > GlobalProtect > Gateways** and click **Satellite Info** in the Info column of the gateway configuration entry. Satellites that have successfully established tunnels with the gateway will display on the **Active Satellites** tab.<br> |
| Step 1 | Verify LSVPN tunnel status on the satellite. | On each firewall hosting a satellite, verify the tunnel status by selecting **Network > IPSec Tunnels** and verify active Status as indicated by a green icon.<br> |

# LSVPN Quick Configs

The following sections provide step-by-step instructions for configuring some common GlobalProtect LSVPN deployments:

- Basic LSVPN Configuration with Static Routing
- Advanced LSVPN Configuration with Dynamic Routing


© Palo Alto Networks, Inc.

# Basic LSVPN Configuration with Static Routing

This quick config shows the fastest way to get up and running with LSVPN. In this example, a single firewall at the corporate headquarters site is configured as both a portal and a gateway. Satellite devices can be quickly and easily deployed with minimal configuration for optimized scalability.

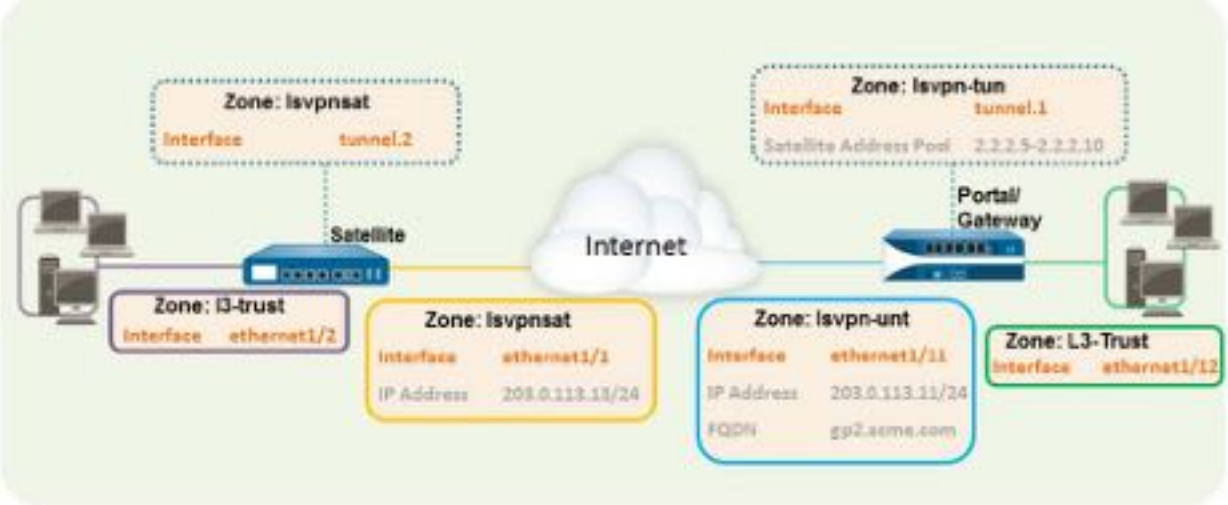


The following workflow shows the steps for setting up this basic configuration:

| Quick Config: Basic LSVPN with Static Routing | | |
|---|---|---|
| Step 1 | Configure a Layer 3 interface. | In this example, the Layer 3 interface on the portal/gateway requires the following configuration:<br>• **Interface**—ethernet1/11<br>• **Security Zone**—lsvpn-unt<br>• **IPv4**—203.0.113.11/24 |
| Step 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.<br>To enable visibility into users and groups connecting over the VPN, enable User-ID in the zone where the VPN tunnels terminate. | In this example, the Tunnel interface on the portal/gateway requires the following configuration:<br>• **Interface**—tunnel.1<br>• **Security Zone**—lsvpn-tun |

© Palo Alto Networks, Inc.

| Quick Config: Basic LSVPN with Static Routing (Continued) | | |
|---|---|---|
| Step 3 | Create the security policy rule to enable traffic flow between the VPN zone where the tunnel terminates (lsvpn-tun) and the trust zone where the corporate applications reside (L3-Trust). | The following is an example of a policy rule for allowing traffic to flow between the VPN zone and the trust zone:<br> 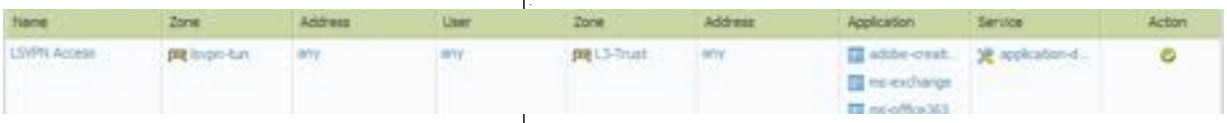 |
| Step 4 | Assign an SSL/TLS Service profile to the portal/gateway. The profile must reference a self-signed server certificate.<br>The certificate subject name must match the FQDN or IP address of the Layer 3 interface you create for the portal/gateway. | 1. On the firewall hosting the GlobalProtect portal, create the root CA certificate for signing the certificates of the GlobalProtect components. In this example, the root CA certificate, `lsvpn-CA`, will be used to issue the server certificate for the portal/gateway. In addition, the portal will use this root CA certificate to sign the CSRs from the satellite devices.<br>2. Create SSL/TLS service profiles for the GlobalProtect portal and gateways.<br>Because the portal and gateway are on the same interface in this example, they can share an SSL/TLS Service profile that uses the same server certificate. In this example, the profile is named `lsvpnserver`. |
| Step 5 | Create a certificate profile. | In this example, the certificate profile `lsvpn-profile`, references the root CA certificate `lsvpn-CA`. The gateway will use this certificate profile to authenticate satellite devices attempting to establish VPN tunnels. |
| Step 6 | Configure an authentication profile for the portal to use if the satellite device serial number is not available. | 1. Create one type of server profile on the portal:<br>• Configure a RADIUS Server Profile.<br>• Configure a TACACS+ Server Profile.<br>• Configure an LDAP Server Profile. If you use LDAP to connect to Active Directory (AD), create a separate LDAP server profile for every AD domain.<br>• Configure a Kerberos Server Profile.<br>2. Configure an authentication profile. In this example, the profile `lsvpn-sat` is used to authenticate satellite devices. |

 © Palo Alto Networks, Inc.

| Quick Config: Basic LSVPN with Static Routing (Continued) | |
|---|---|
| Step 7 Configure the Gateway for LSVPN. | Select **Network > GlobalProtect > Gateways** and **Add** a configuration. This example requires the following gateway configuration:<br>• **Interface**—ethernet1/11<br>• **IP Address**—203.0.113.11/24<br>• **SSL/TLS Server Profile**—lsvpnserver<br>• **Certificate Profile**—lsvpn-profile<br>• **Tunnel Interface**—tunnel.1<br>• **Primary DNS**/**Secondary DNS**—4.2.2.1/4.2.2.2<br>• **IP Pool**—2.2.2.111-2.2.2.120<br>• **Access Route**—10.2.10.0/24 |
| Step 8 Configure the Portal for LSVPN. | Select **Network > GlobalProtect > Portal** and **Add** a configuration. This example requires the following portal configuration:<br>• **Interface**—ethernet1/11<br>• **IP Address**—203.0.113.11/24<br>• **SSL/TLS Server Profile**—lsvpnserver<br>• **Authentication Profile**—lsvpn-sat |
| Step 9 Create a GlobalProtect Satellite Configuration. | On the **Satellite Configuration** tab in the portal configuration, **Add** a Satellite Configuration and a Trusted Root CA and specify the CA the portal will use to issue certificates for the satellites. In this example the required settings are as following:<br>• **Gateway**—203.0.113.11<br>• **Issuing Certificate**—lsvpn-CA<br>• **Trusted Root CA**—lsvpn-CA |

| Quick Config: Basic LSVPN with Static Routing (Continued) | |
|---|---|
| Step 10 Prepare the Satellite Device to Join the LSVPN. | The satellite configuration in this example requires the following settings:<br>Interface Configuration<br>• Layer 3 interface—ethernet1/1, 203.0.113.13/24<br>• Tunnel interface—tunnel.2<br>• Zone—lsvpnsat<br>Root CA Certificate from Portal<br>• lsvpn-CA<br>IPSec Tunnel Configuration<br>• **Tunnel Interface**—tunnel.2<br>• **Portal Address**—203.0.113.11<br>• **Interface**—ethernet1/1<br>• **Local IP Address**—203.0.113.13/24<br>• **Publish all static and connected routes to Gateway**—enabled |

 © Palo Alto Networks, Inc.

# Advanced LSVPN Configuration with Dynamic Routing

In larger LSVPN deployments with multiple gateways and many satellites, investing a little more time in the initial configuration to set up dynamic routing will simplify the maintenance of gateway configurations because access routes will update dynamically. The following example configuration shows how to extend the basic LSVPN configuration to configure OSPF as the dynamic routing protocol.

Setting up an LSVPN to use OSPF for dynamic routing requires the following additional steps on the gateways and the satellites:

- Manual assignment of IP addresses to tunnel interfaces on all gateways and satellites.
- Configuration of OSPF point-to-multipoint (P2MP) on the virtual router on all gateways and satellites. In addition, as part of the OSPF configuration on each gateway, you must manually define the tunnel IP address of each satellite as an OSPF neighbor. Similarly, on each satellite, you must manually define the tunnel IP address of each gateway as an OSPF neighbor.

Although dynamic routing requires additional setup during the initial configuration of the LSVPN, it reduces the maintenance tasks associated with keeping routes up to date as topology changes occur on your network.

The following figure shows an LSVPN dynamic routing configuration. This example shows how to configure OSPF as the dynamic routing protocol for the VPN.

For a basic setup of a LSVPN, follow the steps in Basic LSVPN Configuration with Static Routing. You can then complete the steps in the following workflow to extend the configuration to use dynamic routing rather than static routing.

| Quick Config: LSVPN with Dynamic Routing | |
|---|---|
| Step 1 Add an IP address to the tunnel interface configuration on each gateway and each satellite. | Complete the following steps on each gateway and each satellite:<br>1. Select **Network > Interfaces > Tunnel** and select the tunnel configuration you created for the LSVPN to open the Tunnel Interface dialog.<br>If you have not yet created the tunnel interface, see Step 2 in Quick Config: Basic LSVPN with Static Routing.<br>2. On the **IPv4** tab, click **Add** and then enter an IP address and subnet mask. For example, to add an IP address for the gateway tunnel interface you would enter 2.2.2.100/24.<br>3. Click **OK** to save the configuration. |
| Step 2 Configure the dynamic routing protocol on the gateway.<br> 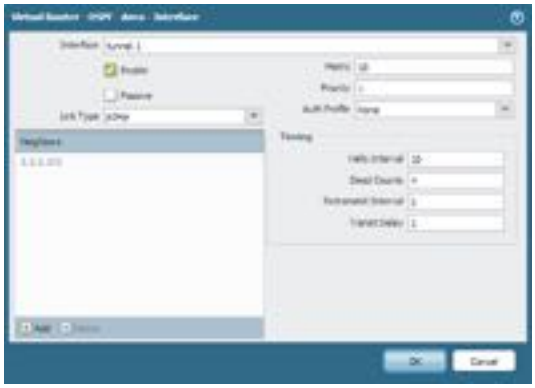 | To configure OSPF on the gateway:<br>1. Select **Network > Virtual Routers** and select the virtual router associated with your VPN interfaces.<br>2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.<br>3. If you are creating a new area, enter an **Area ID** on the **Type** tab.<br>4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.<br>5. Select **p2mp** as the **Link Type**.<br>6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each satellite device, for example 2.2.2.111.<br>7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.<br>8. Repeat this step each time you add a new satellite to the LSVPN. |
| Step 3 Configure the dynamic routing protocol on the satellite.<br> 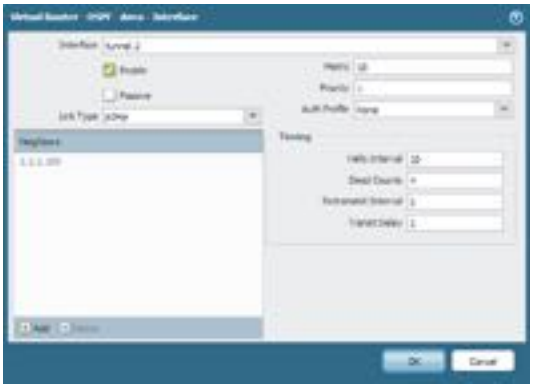 | To configure OSPF on the satellite:<br>1. Select **Network > Virtual Routers** and select the virtual router associated with your VPN interfaces.<br>2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.<br>3. If you are creating a new area, enter an **Area ID** on the **Type** tab.<br>4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.<br>5. Select **p2mp** as the **Link Type**.<br>6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each GlobalProtect gateway, for example 2.2.2.100.<br>7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.<br>8. Repeat this step each time you add a new gateway. |

 © Palo Alto Networks, Inc.

| Quick Config: LSVPN with Dynamic Routing (Continued) | |
|---|---|
| Step 4 Verify that the gateways and satellites are able to form router adjacencies. | • On each satellite and each gateway, confirm that peer adjacencies have formed and that routing table entries have been created for the peers (that is, the satellites have routes to the gateways and the gateways have routes to the satellites). Select **Network > Virtual Router** and click the **More Runtime Stats** link for the virtual router you are using for the LSVPN. On the Routing tab, verify that the LSVPN peer has a route.<br> <br>• On the **OSPF > Interface** tab, verify that the **Type** is **p2mp**.<br> <br>• On the **OSPF > Neighbor** tab, verify that the firewalls hosting your gateways have established router adjacencies with the firewalls hosting your satellites and vice versa. Also verify that the **Status** is **Full**, indicating that full adjacencies have been established.<br> 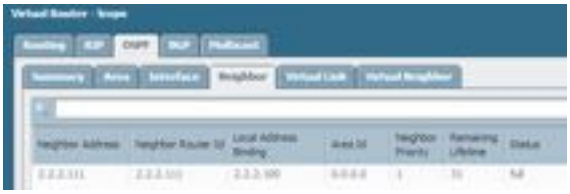 |

 © Palo Alto Networks, Inc.

# Networking

All Palo Alto Networks® next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, and enables you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports. The Interface Deployments section provides basic information on each type of deployment. For more detailed deployment information, refer to Designing Networks with Palo Alto Networks Firewalls.

The following topics describe networking concepts and how to integrate Palo Alto Networks next-generation firewalls into your network.

- Interface Deployments
- Virtual Routers
- Static Routes
- RIP
- OSPF
- BGP
- Session Settings and Timeouts
- DHCP
- NAT
- NPTv6
- LACP
- ECMP
- LLDP

For information on route distribution, refer to Understanding Route Redistribution and Filtering.

# Interface Deployments

A Palo Alto Networks firewall can operate in multiple deployments at once because the deployments occur at the interface level. The following sections describe the deployments.


- Virtual Wire Deployments
- Layer 2 Deployments
- Layer 3 Deployments
- Tap Mode Deployments


## Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together and should be used only when no switching or routing is needed.

A virtual wire deployment allows the following conveniences:

- Simplifies installation and configuration.
- Does not require any configuration changes to surrounding or adjacent network devices.

The virtual wire deployment shipped as the factory default configuration (default-vwire) binds together Ethernet ports 1 and 2 and allows all untagged traffic. You can, however, use a virtual wire to connect any two ports and configure it to block or allow traffic based on the virtual LAN (VLAN) tags; the VLAN tag "0" indicates untagged traffic. You can also create multiple subinterfaces, add them into different zones and then classify traffic according to a VLAN tag, or a combination of a VLAN tag with IP classifiers (address, range, or subnet) to apply granular policy control for specific VLAN tags or for VLAN tags from a specific source IP address, range, or subnet.

**Figure: Virtual Wire Deployment**

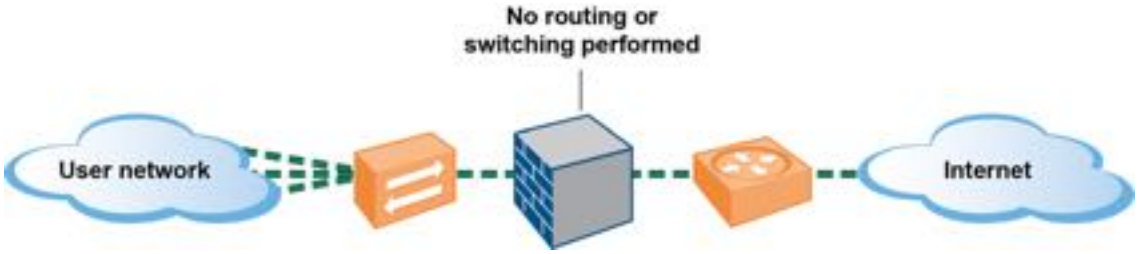


### Virtual Wire Subinterfaces

Virtual wire subinterfaces provide flexibility in enforcing distinct policies when you need to manage traffic from multiple customer networks. It allows you to separate and classify traffic into different zones (the zones can belong to separate virtual systems, if required) using the following criteria:

 © Palo Alto Networks, Inc.

- **VLAN tags** —The example in Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only), shows an Internet Service Provider (ISP) using virtual wire subinterfaces with VLAN tags to separate traffic for two different customers.
- **VLAN tags in conjunction with IP classifiers (address, range, or subnet)**— The following example shows an ISP with two separate virtual systems on a firewall that manages traffic from two different customers. On each virtual system, the example illustrates how virtual wire subinterfaces with VLAN tags and IP classifiers are used to classify traffic into separate zones and apply relevant policy for customers from each network.

| Virtual Wire Subinterface Workflow |
|---|
| 1. Configure two Ethernet interfaces as type virtual wire, and assign these interfaces to a virtual wire. |
| 2. Create subinterfaces on the parent Virtual Wire to separate CustomerA and CustomerB traffic. Make sure that the VLAN tags defined on each pair of subinterfaces that are configured as virtual wire(s) are identical. This is essential because a virtual wire does not switch VLAN tags. |
| 3. Create new subinterfaces and define IP classifiers. This task is optional and only required if you wish to add additional subinterfaces with IP classifiers for further managing traffic from a customer based on the combination of VLAN tags and a specific source IP address, range or subnet.<br>You can also use IP classifiers for managing untagged traffic. To do so, you must create a sub-interface with the vlan tag "0", and define sub-interface(s) with IP classifiers for managing untagged traffic using IP classifiers |

IP classification may only be used on the subinterfaces associated with one side of the virtual wire. The subinterfaces defined on the corresponding side of the virtual wire must use the same VLAN tag, but must not include an IP classifier.

**Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only)**

Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only) depicts CustomerA and CustomerB connected to the firewall through one physical interface, ethernet1/1, configured as a Virtual Wire; it is the ingress interface. A second physical interface, ethernet1/2, is also part of the Virtual Wire; it is the egress interface that provides access to the Internet. For CustomerA, you also have subinterfaces ethernet1/1.1 (ingress) and ethernet1/2.1 (egress). For CustomerB, you have the subinterface ethernet1/1.2 (ingress) and ethernet1/2.2 (egress). When configuring the subinterfaces, you must assign the appropriate VLAN tag and zone in order to apply policies for each customer. In this example, the policies for CustomerA are created between Zone1 and Zone2, and policies for CustomerB are created between Zone3 and Zone4.

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this example, a single subinterface matches the VLAN tag on the incoming packet, hence that subinterface is selected. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

The same VLAN tag must not be defined on the parent virtual wire interface and the subinterface. Verify that the VLAN tags defined on the Tag Allowed list of the parent virtual wire interface (**Network > Virtual Wires**) are not included on a subinterface.

Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers) depicts CustomerA and CustomerB connected to one physical firewall that has two virtual systems (vsys), in addition to the default virtual system (vsys1). Each virtual system is an independent virtual firewall that is managed separately for each customer. Each vsys has attached interfaces/subinterfaces and security zones that are managed independently.

**Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers)**

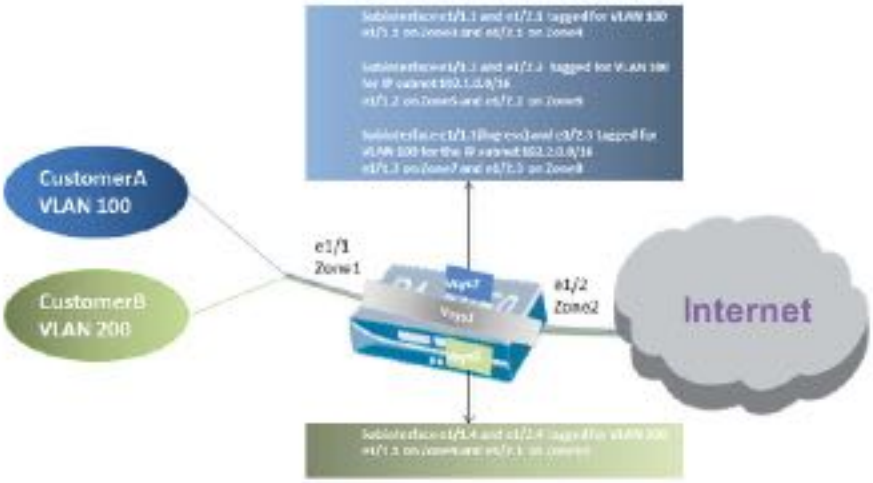


Vsys1 is set up to use the physical interfaces ethernet1/1 and ethernet1/2 as a virtual wire; ethernet1/1 is the ingress interface and ethernet1/2 is the egress interface that provides access to the Internet. This virtual wire is configured to accept all tagged and untagged traffic with the exception of VLAN tags 100 and 200 that are assigned to the subinterfaces.

CustomerA is managed on vsys2 and CustomerB is managed on vsys3. On vsys2 and vsys3, the following vwire subinterfaces are created with the appropriate VLAN tags and zones to enforce policy measures.

| Customer | Vsys | Vwire Subinterfaces | Zone | VLAN Tag | IP Classifier |
|---|---|---|---|---|---|
| A | 2 | e1/1.1 (ingress)<br>e1/2.1 (egress) | Zone3<br>Zone4 | 100<br>100 | None |
| | 2 | e1/1.2 (ingress)<br>e1/2.2 (egress) | Zone5<br>Zone6 | 100<br>100 | IP subnet<br>192.1.0.0/16 |
| | 2 | e1/1.3 (ingress)<br>e1/2.3 (egress) | Zone7<br>Zone8 | 100<br>100 | IP subnet<br>192.2.0.0/16 |
| B | 3 | e1/1.4 (ingress)<br>e1/2.4 (egress) | Zone9<br>Zone10 | 200<br>200 | None |

 © Palo Alto Networks, Inc.

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this case, for CustomerA, there are multiple subinterfaces that use the same VLAN tag. Hence, the firewall first narrows the classification to a subinterface based on the source IP address in the packet. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

For return-path traffic, the firewall compares the destination IP address as defined in the IP classifier on the customer-facing subinterface and selects the appropriate virtual wire to route traffic through the accurate subinterface.

The same VLAN tag must not be defined on the parent virtual wire interface and the subinterface. Verify that the VLAN tags defined on the Tag Allowed list of the parent virtual wire interface (**Network > Virtual Wires**) are not included on a subinterface.

## Layer 2 Deployments

In a Layer 2 deployment, the firewall provides switching between two or more networks. Each group of interfaces must be assigned to a VLAN object in order for the firewall to switch between them. The firewall will perform VLAN tag switching when layer 2 subinterfaces are attached to a common VLAN object. Choose this option when switching is required.

**Figure: Layer 2 Deployment**

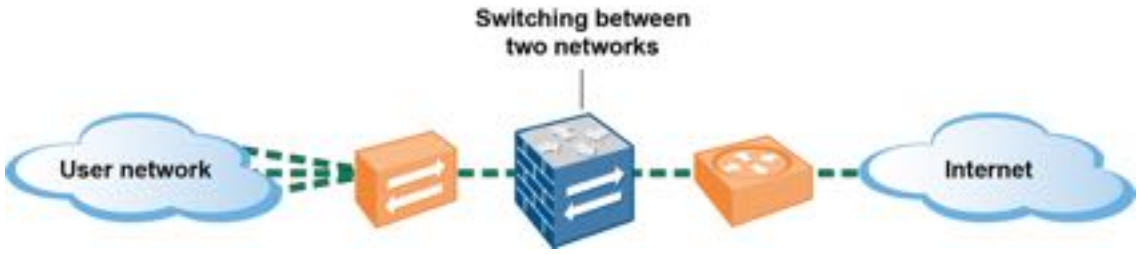


## Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between multiple ports. An IP address must be assigned to each interface and a virtual router must be defined to route the traffic. Choose this option when routing is required.

**Figure: Layer 3 Deployment**

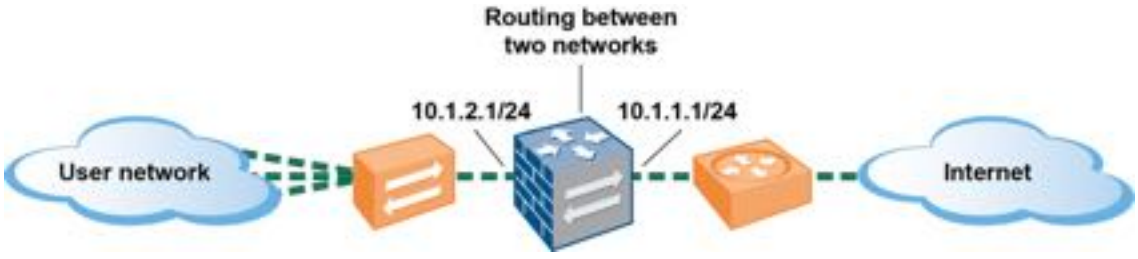

In addition, because the firewall must route traffic in a Layer 3 deployment, you must configure a virtual router. See Virtual Routers.

### Point-to-Point Protocol over Ethernet Support

You can configure the firewall to be a Point-to-Point Protocol over Ethernet (PPPoE) termination point to support connectivity in a Digital Subscriber Line (DSL) environment where there is a DSL modem but no other PPPoE device to terminate the connection.

You can choose the PPPoE option and configure the associated settings when an interface is defined as a Layer 3 interface.

PPPoE is not supported in HA active/active mode.

### DHCP Client

You can configure the firewall interface to act as a DHCP client and receive a dynamically assigned IP address. The firewall also provides the capability to propagate settings received by the DHCP client interface into a DHCP server operating on the firewall. This is most commonly used to propagate DNS server settings from an Internet service provider to client machines operating on the network protected by the firewall.

DHCP client is not supported in HA active/active mode.

For more information, see DHCP.

## Tap Mode Deployments

A network tap is a device that provides a way to access data flowing across a computer network. Tap mode deployment allows you to passively monitor traffic flows across a network by way of a switch SPAN or mirror port.

The SPAN or mirror port permits the copying of traffic from other ports on the switch. By dedicating an interface on the firewall as a tap mode interface and connecting it with a switch SPAN port, the switch SPAN port provides the firewall with the mirrored traffic. This provides application visibility within the network without being in the flow of network traffic.

When deployed in tap mode, the firewall is not able to take action, such as block traffic or apply QoS traffic control.

 © Palo Alto Networks, Inc.

# Virtual Routers

The firewall uses virtual routers to obtain routes to other subnets by manually defining a route (static routes) or through participation in Layer 3 routing protocols (dynamic routes). The best routes obtained through these methods are used to populate the firewall's IP route table. When a packet is destined for a different subnet, the Virtual Router obtains the best route from this IP route table and forwards the packet to the next hop router defined in the table.

The Ethernet interfaces and VLAN interfaces defined on the firewall receive and forward the Layer 3 traffic. The destination zone is derived from the outgoing interface based on the forwarding criteria, and policy rules are consulted to identify the security policies to be applied. In addition to routing to other network devices, virtual routers can route to other virtual routers within the same firewall if a next hop is specified to point to another virtual router.

You can configure the virtual router to participate with dynamic routing protocols (BGP, OSPF, or RIP) as well as adding static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that are not shared between virtual routers, enabling you to configure different routing behaviors for different interfaces.

Each Layer 3 interface, loopback interface, and VLAN interface defined on the firewall must be associated with a virtual router. While each interface can belong to only one virtual router, multiple routing protocols and static routes can be configured for a virtual router. Regardless of the static routes and dynamic routing protocols configured for a virtual router, a common general configuration is required. The firewall uses Ethernet switching to reach other devices on the same IP subnet.

The following Layer 3 routing protocols are supported from Virtual Routers:

- RIP
- OSPF
- OSPFv3
- BGP

| **Define a Virtual Router General Configuration** | | |
|---|---|---|
| Step 1 | Gather the required information from your network administrator. | • Interfaces that you want to route<br>• Administrative distances for static, OSFP internal, OSPF external, IBGP, EBGP and RIP |
| Step 2 | Create the virtual router and name it. | 1. Select **Network > Virtual Routers**.<br>2. Click **Add** and enter a name for the virtual router.<br>3. Select interfaces to apply to the virtual router.<br>4. Click **OK**. |
| Step 3 | Select interfaces to apply to the virtual router. | 1. Click **Add** in the Interfaces box.<br>2. Select an already defined interface from the drop-down.<br>3. Repeat Step 2 for all interfaces that you want to add to the virtual router. |

© Palo Alto Networks, Inc.

| Define a Virtual Router General Configuration (Continued) | | |
|---|---|---|
| Step 4 | Set Administrative Distances for static and dynamic routing. | Set Administrative Distances as required.<br>• **Static**—Range is 10-240; default is 10.<br>• **OSPF Internal**—Range is 10-240; default is 30.<br>• **OSPF External**—Range is 10-240; default is 110.<br>• **IBGP**—Range is 10-240; default is 200.<br>• **EBGP**—Range is 10-240; default is 20.<br>• **RIP**—Range is 10-240; default is 120. |
| Step 5 | Save virtual router general settings. | Click **OK** to save your settings. |
| Step 6 | Commit your changes. | Click **Commit**. The device may take up to 90 seconds to save your changes. |

 © Palo Alto Networks, Inc.

# Static Routes

The following procedure shows how to integrate the firewall into the network using static routing.

| Set Up Interfaces and Zones | | |
|---|---|---|
| Step 1 | Configure a default route to your Internet router. | 1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog.<br>2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and enter the route in the **Destination** field (for example, 0.0.0.0/0).<br>3. Select the **IP Address** radio button in the **Next Hop** field and then enter the IP address and netmask for your Internet gateway (for example, 208.80.56.1).<br>4. Click **OK** twice to save the virtual router configuration. |
| Step 2 | Configure the external interface (the interface that connects to the Internet). | 1. Select **Network > Interfaces** and then select the interface you want to configure. In this example, we are configuring Ethernet1/3 as the external interface.<br>2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.<br>3. In the **Virtual Router** drop-down, select **default**.<br>4. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Untrust, and then click **OK**.<br>5. To assign an IP address to the interface, select the **IPv4** tab and **Static** radio button. Click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 208.80.56.100/24.<br>6. To enable you to ping the interface, select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. Enter a **Name** for the profile, select **Ping** and then click **OK**.<br>7. To save the interface configuration, click **OK**. |

| Set Up Interfaces and Zones (Continued) | | |
|---|---|---|
| Step 3 | Configure the interface that connects to your internal network.<br><br>In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT. See Configure NAT for details. | 1. Select **Network > Interfaces** and select the interface you want to configure. In this example, we are configuring Ethernet1/4 as the internal interface.<br>2. Select **Layer3** from the **Interface Type** drop-down.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Trust, and then click **OK**.<br>4. Select the same Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab and the **Static** radio button, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 4 | Configure the interface that connects to the DMZ. | 1. Select the interface you want to configure.<br>2. Select **Layer3** from the **Interface Type** drop-down. In this example, we are configuring Ethernet1/13 as the DMZ interface.<br>3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example DMZ, and then click **OK**.<br>4. Select the Virtual Router you used in Step 2, default in this example.<br>5. To assign an IP address to the interface, select the **IPv4** tab and the **Static** radio button, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24.<br>6. To enable you to ping the interface, select the management profile that you created in Step 2-6.<br>7. To save the interface configuration, click **OK**. |
| Step 5 | Save the interface configuration. | Click **Commit**. |
| Step 6 | Cable the firewall. | Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment. |
| Step 7 | Verify that the interfaces are active. | From the web interface, select **Network > Interfaces** and verify that icon in the Link State column is green. You can also monitor link state from the **Interfaces** widget on the **Dashboard**.<br><br> 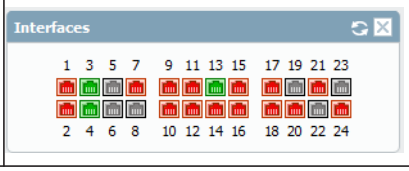 |

 © Palo Alto Networks, Inc.

# RIP

Routing Information Protocol (RIP) is an interior gateway protocol (IGP) that was designed for small IP networks. RIP relies on hop count to determine routes; the best routes have the fewest number of hops. RIP is based on UDP and uses port 520 for route updates. By limiting routes to a maximum of 15 hops, the protocol helps prevent the development of routing loops, but also limits the supported network size. If more than 15 hops are required, traffic is not routed. RIP also can take longer to converge than OSPF and other routing protocols. The firewall supports RIP v2.

Perform the following procedure to configure RIP.

| **Configure RIP** | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Virtual Routers for details. |
| Step 2 | Configure general RIP configuration settings. | 1. Select the **RIP** tab.<br>2. Select the **Enable** check box to enable the RIP protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through RIP. This is the recommended default setting.<br>4. De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through RIP. |
| Step 3 | Configure interfaces for the RIP protocol. | 1. On **Interfaces** tab, select an interface from the drop-down in the Interface configuration section.<br>2. Select an already defined interface from the drop-down.<br>3. Select the **Enable** check box.<br>4. Select the **Advertise** check box to advertise a default route to RIP peers with the specified metric value.<br>5. You can optionally select a profile from the **Auth Profile** drop-down. See Step 5 for details.<br>6. Select normal, passive or send-only from the **Mode** drop-down.<br>7. Click **OK**. |
| Step 4 | Configure RIP timers. | 1. On the **Timers** tab, enter a value in the **Interval Seconds (sec)** box. This setting defines the length of the timer interval in seconds (range is 1-60; default is 1). This duration is used for the remaining RIP timing fields.<br>2. Enter an **Update Intervals** value to define the number of intervals between route update announcements (range is 1-3600; default is 30).<br>3. Specify the **Delete Intervals** to define the number of intervals between the time that the route expires to its deletion (range is 1-3600; default is 180).<br>4. Specify the **Expire Intervals** to define the number of intervals between the time that the route was last updated to its expiration (range is 1-3600; default is 120). |

© Palo Alto Networks, Inc.

| Configure RIP | |
|---|---|
| Step 5 (Optional) Configure Auth Profiles. | By default, the firewall does not use RIP authentication for the exchange between RIP neighbors. Optionally, you can configure RIP authentication between RIP neighbors by either a simple password or using MD5 authentication.<br>**Simple Password RIP authentication**<br>1. Select **Auth Profiles** and click **Add**.<br>2. Enter a name for the authentication profile to authenticate RIP messages.<br>3. Select **Simple Password** as the **Password Type**.<br>4. Enter a simple password and then confirm.<br>**MD5 RIP authentication**<br>1. Select **Auth Profiles** and click **Add**.<br>2. Enter a name for the authentication profile to authenticate RIP messages.<br>3. Select **MD5** as the **Password Type**.<br>4. Click **Add**.<br>5. Enter one or more password entries, including:<br>• Key-ID (range is 0-255)<br>• Key<br>6. You can optionally select **Preferred** status.<br>7. Click **OK** to specify the key to be used to authenticate outgoing message.<br>8. Click **OK** again in the Virtual Router - RIP Auth Profile dialog box. |

 © Palo Alto Networks, Inc.

# OSPF

Open Shortest Path First (OSPF) is an interior gateway protocol (IGP) that is most often used to dynamically manage network routes in large enterprise network. It determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The information gathered from the LSAs is used to construct a topology map of the network. This topology map is shared across routers in the network and used to populate the IP routing table with available routes.

Changes in the network topology are detected dynamically and used to generate a new topology map within seconds. A shortest path tree is computed of each route. Metrics associated with each routing interface are used to calculate the best route. These can include distance, network throughput, link availability etc. Additionally, these metrics can be configured statically to direct the outcome of the OSPF topology map.

Palo Alto networks implementation of OSPF fully supports the following RFCs:

- RFC 2328 (for IPv4)
- RFC 5340 (for IPv6)

The following topics provide more information about the OSPF and procedures for configuring OSPF on the firewall:

- OSPF Concepts
- Configure OSPF
- Configure OSPFv3
- Configure OSPF Graceful Restart
- Confirm OSPF Operation

Also refer to How to Configure OSPF Tech Note.

## OSPF Concepts

The following topics introduce the OSPF concepts you will need to understand in order to configure the firewall to participate in an OSPF network:

- OSPFv3
- OSPF Neighbors
- OSPF Areas
- OSPF Router Types

## OSPFv3

OSPFv3 provides support for the OSPF routing protocol within an IPv6 network. As such, it provides support for IPv6 addresses and prefixes. It retains most of the structure and functions in OSPFv2 (for IPv4) with some minor changes. The following are some of the additions and changes to OSPFv3:

- **Support for multiple instances per link**—With OSPFv3, you can run multiple instances of the OSPF protocol over a single link. This is accomplished by assigning an OSPFv3 instance ID number. An interface that is assigned to an instance ID drops packets that contain a different ID.
- **Protocol Processing Per-link**—OSPFv3 operates per-link instead of per-IP-subnet as on OSPFv2.
- **Changes to Addressing**—IPv6 addresses are not present in OSPFv3 packets, except for LSA payloads within link state update packets. Neighboring routers are identified by the Router ID.
- **Authentication Changes**—OSPFv3 doesn't include any authentication capabilities. Configuring OSPFv3 on a firewall requires an authentication profile that specifies Encapsulating Security Payload (ESP) or IPv6 Authentication Header (AH).The re-keying procedure specified in RFC 4552 is not supported in this release.
- **Support for multiple instances per-link**—Each instance corresponds to an instance ID contained in the OSPFv3 packet header.
- **New LSA Types**—OSPFv3 supports two new LSA types: Link LSA and Intra Area Prefix LSA.

All additional changes are described in detail in RFC 5340.

## OSPF Neighbors

Two OSPF-enabled routers connected by a common network and in the same OSPF area that form a relationship are OSPF neighbors. The connection between these routers can be through a common broadcast domain or by a point-to-point connection. This connection is made through the exchange of hello OSPF protocol packets. These neighbor relationships are used to exchange routing updates between routers.

 © Palo Alto Networks, Inc.

## OSPF Areas

OSPF operates within a single autonomous system (AS). Networks within this single AS, however, can be divided into a number of areas. By default, Area 0 is created. Area 0 can either function alone or act as the OSPF backbone for a larger number of areas. Each OSPF area is named using a 32-bit identifier which in most cases is written in the same dotted-decimal notation as an IP4 address. For example, Area 0 is usually written as 0.0.0.0.

The topology of an area is maintained in its own link state database and is hidden from other areas, which reduces the amount of traffic routing required by OSPF. The topology is then shared in a summarized form between areas by a connecting router.

**Table: OSPF Area Types**

| Area Type | Description |
|---|---|
| Backbone Area | The backbone area (Area 0) is the core of an OSPF network. All other areas are connected to it and all traffic between areas must traverse it. All routing between areas is distributed through the backbone area. While all other OSPF areas must connect to the backbone area, this connection doesn't need to be direct and can be made through a virtual link. |
| Normal OSPF Area | In a normal OSPF area there are no restrictions; the area can carry all types of routes. |
| Stub OSPF Area | A stub area does not receive routes from other autonomous systems. Routing from the stub area is performed through the default route to the backbone area. |
| NSSA Area | The Not So Stubby Area (NSSA) is a type of stub area that can import external routes, with some limited exceptions. |

## OSPF Router Types

Within an OSPF area, routers are divided into the following categories.

- **Internal Router**—A router with that has OSPF neighbor relationships only with devices in the same area.
- **Area Border Router (ABR)**—A router that has OSPF neighbor relationships with devices in multiple areas. ABRs gather topology information from their attached areas and distribute it to the backbone area.
- **Backbone Router**—A backbone router is any OSPF router that is attached to the OSPF backbone. Since ABRs are always connected to the backbone, they are always classified as backbone routers.
- **Autonomous System Boundary Router (ASBR)**—An ASBR is a router that attaches to more than one routing protocol and exchanges routing information between them.

© Palo Alto Networks, Inc.

## Configure OSPF

OSPF determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The router keeps information about the links between it and the destination and can make highly efficient routing decisions. A cost is assigned to each router interface, and the best routes are determined to be those with the lowest costs, when summed over all the encountered outbound router interfaces and the interface receiving the LSA.

Hierarchical techniques are used to limit the number of routes that must be advertised and the associated LSAs. Because OSPF dynamically processes a considerable amount of route information, it has greater processor and memory requirements than does RIP.

| **Configure OSPF** | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Virtual Routers for details. |
| Step 2 | Configure general OSPF configuration settings. | 1. Select the **OSPF** tab.<br>2. Select the **Enable** check box to enable the OSPF protocol.<br>3. Optionally enter the **Router ID**.<br>4. Select the **Reject Default Route** check box if you do not want to learn any default routes through OSPF. This is the recommended default setting.<br>De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through OSPF. |


© Palo Alto Networks, Inc.

**Configure OSPF (Continued)**

| | |
|---|---|
| Step 3 Configure Areas - Type for the OSFP protocol. | 1. On the **Areas** tab, click **Add**.<br>2. Enter an Area ID for the area in *x.x.x.x* format. This is the identifier that each neighbor must accept to be part of the same area.<br>3. On the **Type** tab, select one of the following from the area **Type** drop-down:<br>• **Normal**—There are no restrictions; the area can carry all types of routes.<br>• **Stub**—There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, configure the following:<br>– **Accept Summary**—Link state advertisements (LSA) are accepted from other areas. If this option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.<br>– **Advertise Default Route**—Default route LSAs will be included in advertisements to the stub area along with a configured metric value in the configured range 1-255.<br>• **NSSA** (Not-So-Stubby Area)—The firewall can only leave the area by routes other than OSPF routes. If selected, configure **Accept Summary** and **Advertise Default Route** as described for **Stub**. If you select this option, configure the following:<br>– **Type**—Select either **Ext 1** or **Ext 2** route type to advertise the default LSA.<br>– **Ext Ranges**—Click **Add** in the section to enter ranges of external routes that you want to enable or suppress advertising for.<br>4. **Priority**—Enter the OSPF priority for this interface (0-255). This is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR) according to the OSPF protocol. When the value is zero, the router will not be elected as a DR or BDR.<br>• **Auth Profile**—Select a previously-defined authentication profile.<br>• **Timing**—It is recommended that you keep the default timing settings.<br>• **Neighbors**—For p2pmp interfaces, enter the neighbor IP address for all neighbors that are reachable through this interface.<br>5. Select **normal**, **passive** or **send-only** as the **Mode**.<br>6. Click **OK**. |

| Configure OSPF (Continued) | | |
|---|---|---|
| Step 4 | Configure Areas - Range for the OSFP protocol | 1. On the **Range** tab, click **Add** to aggregate LSA destination addresses in the area into subnets.<br>2. **Advertise** or **Suppress** advertising LSAs that match the subnet, and click **OK**. Repeat to add additional ranges. |
| Step 5 | Configure Areas - Interfaces for the OSFP protocol | 1. On the **Interface** tab, click **Add** and enter the following information for each interface to be included in the area:<br>• **Interface**—Select an interface from the drop-down box.<br>• **Enable**—Selecting this option causes the OSPF interface settings to take effect.<br>• **Passive**—Select the check box to if you do not want the OSPF interface to send or receive OSPF packets. Although OSPF packets are not sent or received if you choose this option, the interface is included in the LSA database.<br>• **Link type**—Choose **Broadcast** if you want all neighbors that are accessible through the interface to be discovered automatically by multicasting OSPF hello messages, such as an Ethernet interface. Choose **p2p** (point-to-point) to automatically discover the neighbor. Choose **p2mp** (point-to-multipoint) when neighbors must be defined manually. Defining neighbors manually is allowed only for **p2mp** mode.<br>• **Metric**—Enter an OSPF metric for this interface (range is 0-65535; default is 10).<br>• **Priority**—Enter an OSPF priority for this interface. This is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR) (range is 0-255; default is 1). If zero is configured, the router will not be elected as a DR or BDR.<br>• **Auth Profile**—Select a previously-defined authentication profile.<br>• **Timing**—The following OSPF timing settings can be set. Palo Alto Networks recommends that you retain the default timing settings.<br>– **Hello Interval (sec)**—Interval (in seconds) at which the OSPF process sends hello packets to its directly connected neighbors (range is 0-3600; default is 10).<br>– **Dead Counts—**Number of times the hello interval can occur for a neighbor without OSPF receiving a hello packet from the neighbor, before OSPF considers that neighbor down (range is 3-20; default is 4). The **Hello Interval** multiplied by the **Dead Counts** equals the value of the dead timer. |


© Palo Alto Networks, Inc.

| Configure OSPF (Continued) | |
|---|---|
| | – **Retransmit Interval (sec)**—Length of time (in seconds) that OSPF waits to receive a link state advertisement (LSA) from a neighbor before OSPF retransmits the LSA (range is 0-3600; default is 10).<br>– **Transit Delay (sec)**—Length of time (in seconds) that an LSA is delayed before it is sent out of an interface (range is 0-3600; default is 1).<br>– **Graceful Restart Hello Delay (sec)**—Applies to an OSPF interface when Active/Passive High Availability is configured. **Graceful Restart Hello Delay** is the length of time (in seconds) during which the firewall sends Grace LSA packets at 1-second intervals (range is 1-10; default is 10). During this time, no hello packets are sent from the restarting firewall. During the restart, the dead timer (which is the **Hello Interval** multiplied by the **Dead Counts**) is also counting down. If the dead timer is too short, the adjacency will go down during the graceful restart because of the hello delay. Therefore, it is recommended that the dead timer be at least four times the value of the **Graceful Restart Hello Delay**. For example, a **Hello Interval** of 10 seconds and a **Dead Counts** of 4 yield a dead timer of 40 seconds. If the **Graceful Restart Hello Delay** is set to 10 seconds, that 10-second delay of hello packets is comfortably within the 40-second dead timer, so the adjacency will not time out during a graceful restart.<br>• If **p2mp** is selected for **Link Type** interfaces, enter the neighbor IP addresses for all neighbors that are reachable through this interface.<br>2. Click **OK** |
| Step 6 Configure Areas - Virtual Links. | 1. On the **Virtual Link** tab, click **Add** and enter the following information for each virtual link to be included in the backbone area:<br>• **Name**—Enter a name for the virtual link.<br>• **Neighbor ID**—Enter the router ID of the router (neighbor) on the other side of the virtual link.<br>• **Transit Area**—Enter the area ID of the transit area that physically contains the virtual link.<br>• **Enable**—Select to enable the virtual link.<br>• **Timing**—It is recommended that you keep the default timing settings.<br>• **Auth Profile**—Select a previously-defined authentication profile.<br>2. Click **OK**. |

| Configure OSPF (Continued) | |
|---|---|
| Step 7 (Optional) Configure Auth Profiles. | By default, the firewall does not use OSPF authentication for the exchange between OSPF neighbors. Optionally, you can configure OSPF authentication between OSPF neighbors by either a simple password or using MD5 authentication.<br>**Simple Password OSPF authentication**<br>1. On the **Auth Profiles** tab, click **Add**.<br>2. Enter a name for the authentication profile to authenticate OSPF messages.<br>3. Select **Simple Password** as the **Password Type**.<br>4. Enter a simple password and then confirm.<br>**MD5 OSPF authentication**<br>1. On the **Auth Profiles** tab, click **Add**.<br>2. Enter a name for the authentication profile to authenticate OSPF messages.<br>3. Select **MD5** as the **Password Type**.<br>4. Click **Add**.<br>5. Enter one or more password entries, including:<br>• Key-ID (range is 0-255)<br>• Key<br>• Select the **Preferred** option to specify that the key be used to authenticate outgoing messages.<br>6. Click **OK**.<br>7. Click **OK** again in the Virtual Router - OSPF Auth Profile dialog box. |
| Step 8 Configure Advanced OSPF options. | 1. On the **Advanced** tab, select the **RFC 1583 Compatibility** check box to ensure compatibility with RFC 1583.<br>2. Configure a value for the **SPF Calculation Delay** (sec) timer.<br>This timer allows you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times.<br>3. Configure a value for the **LSA Interval (sec)** time. This timer specifies the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur. |

 © Palo Alto Networks, Inc.

# Configure OSPFv3

| Configure OSPFv3 | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Virtual Routers for details. |
| Step 2 | Configure general OSPF configuration settings. | 1. Select the **OSPF** tab.<br>2. Select the **Enable** check box to enable the OSPF protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through OSPF. This is the recommended default setting.<br>4. Clear the **Reject Default Route** check box if you want to permit redistribution of default routes through OSPF. |
| Step 3 | Configure general OSPFv3 configuration settings. | 1. Select the **OSPFv3** tab.<br>2. Select the **Enable** check box to enable the OSPF protocol.<br>3. Select the **Reject Default Route** check box if you do not want to learn any default routes through OSPFv3 This is the recommended default setting.<br>De-select the **Reject Default Route** check box if you want to permit redistribution of default routes through OSPFv3. |
| Step 4 | Configure Auth Profile for the OSFPv3 protocol.<br>While OSPFv3 doesn't include any authentication capabilities of its own, it relies entirely on IPsec to secure communications between neighbors. | When configuring an authentication profile, you must use Encapsulating Security Payload (ESP) or IPv6 Authentication Header (AH).<br>**ESP OSPFv3 authentication**<br>1. On the **Auth Profiles** tab, click **Add**.<br>2. Enter a name for the authentication profile to authenticate OSPFv3 messages.<br>3. Specify a Security Policy Index (**SPI**). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a hexadecimal value between 00000000 and FFFFFFFF.<br>4. Select **ESP** for **Protocol**.<br>5. Select a **Crypto Algorithm** from the drop-down.<br>You can enter none or one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5.<br>6. If a **Crypto Algorithm** other than none was selected, enter a value for **Key** and then confirm. |

| Configure OSPFv3 (Continued) | |
|---|---|
| | **AH OSPFv3 authentication**<br>1. On the **Auth Profiles** tab, click **Add**.<br>2. Enter a name for the authentication profile to authenticate OSPFv3 messages.<br>3. Specify a Security Policy Index (**SPI**). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a hexadecimal value between 00000000 and FFFFFFFF.<br>4. Select **AH** for **Protocol**.<br>5. Select a **Crypto Algorithm** from the drop-down.<br>You must enter one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5.<br>6. Enter a value for **Key** and then confirm.<br>7. Click **OK**.<br>8. Click **OK** again in the Virtual Router - OSPF Auth Profile dialog. |
| Step 5 Configure Areas - Type for the OSFP protocol. | 1. On the **Areas** tab, click **Add**.<br>2. Enter an Area ID. This is the identifier that each neighbor must accept to be part of the same area.<br>3. On the **General** tab, select one of the following from the area **Type** drop-down:<br>• **Normal**—There are no restrictions; the area can carry all types of routes.<br>• **Stub**—There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, configure the following:<br>– **Accept Summary**—Link state advertisements (LSA) are accepted from other areas. If this option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.<br>– **Advertise Default Route**—Default route LSAs will be included in advertisements to the stub area along with a configured metric value in the configured range 1-255. |

 © Palo Alto Networks, Inc.

| Configure OSPFv3 (Continued) | | |
|---|---|---|
| | | • **NSSA** (Not-So-Stubby Area)—The firewall can only leave the area by routes other than OSPF routes. If selected, configure **Accept Summary** and **Advertise Default Route** as described for **Stub**. If you select this option, configure the following:<br>– **Type**—Select either **Ext 1** or **Ext 2** route type to advertise the default LSA.<br>– **Ext Ranges**—Click **Add** in the section to enter ranges of external routes that you want to enable or suppress advertising for. |
| Step 6 | Associate an OSPFv3 authentication profile to an area or an interface. | **To an Area**<br>1. On the **Areas** tab, select an existing area from the table.<br>2. On the **General** tab, select a previously defined **Authentication Profile** from the **Authentication** drop-down.<br>3. Click **OK**.<br>**To an Interface**<br>1. On the **Areas** tab, select an existing area from the table.<br>2. Select the **Interface** tab and click **Add**.<br>3. Select the authentication profile you want to associate with the OSPF interface from the **Auth Profile** drop-down. |
| Step 7 | (Optional) Configure Export Rules | 1. On the **Export** tab, click **Add**.<br>2. Select the **Allow Redistribute Default Route** check box to permit redistribution of default routes through OSPFv3.<br>3. Select the name of a redistribution profile. The value must be an IP subnet or valid redistribution profile name.<br>4. Select a metric to apply for **New Path Type**.<br>5. Specify a **New Tag** for the matched route that has a 32-bit value.<br>6. Assign a metric for the new rule (range is 1 - 65535).<br>7. Click **OK**. |

| Configure OSPFv3 (Continued) | |
|---|---|
| Step 8 Configure Advanced OSPFv3 options. | 1. On the **Advanced** tab, select the **Disable Transit Routing for SPF Calculation** check box if you want the firewall to participate in OSPF topology distribution without being used to forward transit traffic.<br>2. Configure a value for the **SPF Calculation Delay** (sec) timer.<br>This timer allows you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times.<br>3. Configure a value for the **LSA Interval (sec) time**, This timer specifies the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur.<br>4. Optionally, Configure OSPF Graceful Restart. |

## Configure OSPF Graceful Restart

OSPF Graceful Restart directs OSPF neighbors to continue using routes through a device during a short transition when it is out of service. This behavior increases network stability by reducing the frequency of routing table reconfiguration and the related route flapping that can occur during short periodic down times.

For a Palo Alto Networks firewall, OSPF Graceful Restart involves the following operations:

- **Firewall as a restarting device**—In a situation where the firewall will be down for a short period of time or is unavailable for short intervals, it sends Grace LSAs to its OSPF neighbors. The neighbors must be configured to run in Graceful Restart Helper mode. In Helper Mode, the neighbors receive the Grace LSAs that inform it that the firewall will perform a graceful restart within a specified period of time defined as the **Grace Period**. During the grace period, the neighbor continues to forward routes through the firewall and to send LSAs that announce routes through the firewall. If the firewall resumes operation before expiration of the grace period, traffic forwarding will continue as before without network disruption. If the firewall does not resume operation after the grace period has expired, the neighbors will exit helper mode and resume normal operation, which will involve reconfiguring the routing table to bypass the firewall.
- **Firewall as a Graceful Restart Helper**—In a situation where neighboring routers may be down for a short periods of time, the firewall can be configured to operate in Graceful Restart Helper mode. If configured in this mode, the firewall will be configured with a **Max Neighbor Restart Time**. When the firewall receives the Grace LSAs from its OSFP neighbor, it will continue to route traffic to the neighbor and advertise routes through the neighbor until either the grace period or max neighbor restart time expires. If neither expires before the neighbor returns to service, traffic forwarding continues as before without network disruption. If either period expires before the neighbor returns to service, the firewall will exit helper mode and resume normal operation, which will involve reconfiguring the routing table to bypass the neighbor.

 © Palo Alto Networks, Inc.

| Configure OSPF Graceful Restart |
|---|
| 1. Select **Network > Virtual Routers** and select the virtual router you want to configure. |
| 2. Select **OSPF > Advanced**. |
| 3. Verify that the following check boxes are selected (they are enabled by default):<br>• **Enable Graceful Restart**<br>• **Enable Helper Mode**<br>• **Enable Strict LSA checking**<br>These check boxes should remain selected unless required by your topology. |
| 4. Configure a **Grace Period** in seconds. |
| 5. Configure a **Max Neighbor Restart Time** in seconds. |

## Confirm OSPF Operation

Once an OSPF configuration has been committed, you can use any of the following operations to confirm that OSPF is operating:

- View the Routing Table
- Confirm OSPF Adjacencies
- Confirm that OSPF Connections are Established

### View the Routing Table

By viewing the routing table, you can see whether OSPF routes have been established. The routing table is accessible from either the web interface or the CLI. If you are using the CLI, use the following commands:

- `show routing route`
- `show routing fib`

The following procedure describes how to use the web interface to view the routing table.

| View the Routing Table |
|---|
| 1. Select **Network > Virtual Routers**. |
| 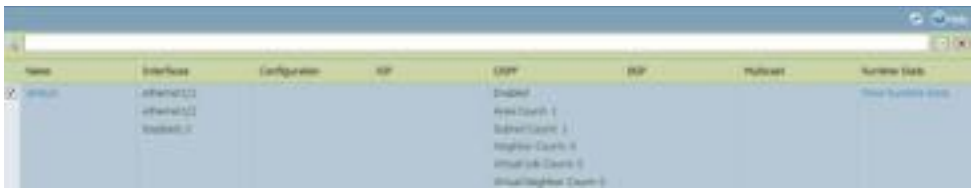 |

**View the Routing Table**

2. Select the **Routing** tab and examine the **Flags** column of the routing table for routes that were learned by OSPF.

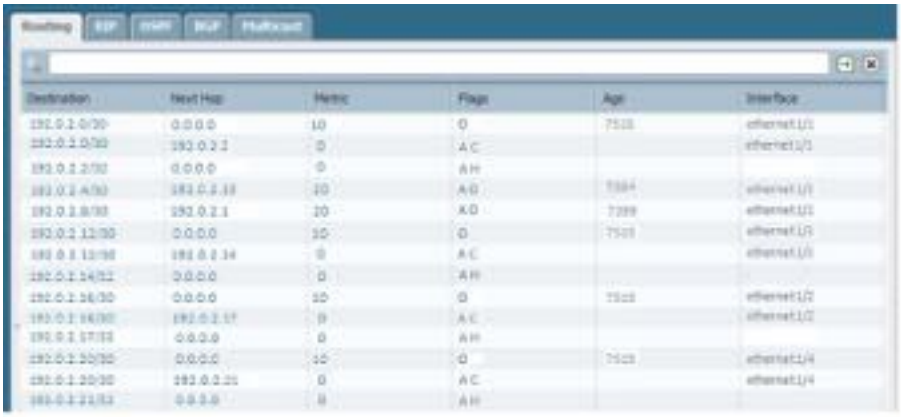

## Confirm OSPF Adjacencies

By viewing the **Neighbor** tab as described in the following procedure, you can confirm that OSPF adjacencies have been established.

**View the Neighbor Tab to Confirm OSPF Adjacencies**

1. Select **Network > Virtual Routers**.

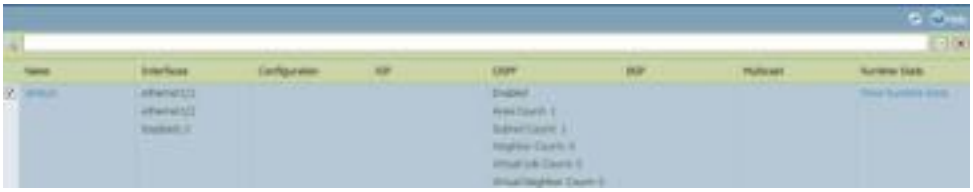

2. Select **OSPF > Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established.

 © Palo Alto Networks, Inc.

## Confirm that OSPF Connections are Established

By viewing the system log, you can confirm that OSPF connections have been established, as described in the following procedure:

**Examine the System Log**

1. Select **Monitor> System** and look for messages to confirm that OSPF adjacencies have been established.

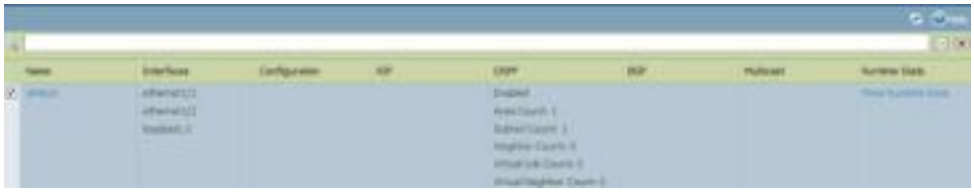

2. Select **OSPF > Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established.

# BGP

Border Gateway Protocol (BGP) is the primary Internet routing protocol. BGP determines network reachability based on IP prefixes that are available within autonomous systems (AS), where an AS is a set of IP prefixes that a network provider has designated to be part of a single routing policy.

In the routing process, connections are established between BGP peers (or neighbors). If a route is permitted by the policy, it is stored in the routing information base (RIB). Each time the local firewall RIB is updated, the firewall determines the optimal routes and sends an update to the external RIB, if export is enabled.

Conditional advertisement is used to control how BGP routes are advertised. The BGP routes must satisfy conditional advertisement rules before being advertised to peers.

BGP supports the specification of aggregates, which combine multiple routes into a single route. During the aggregation process, the first step is to find the corresponding aggregation rule by performing a longest match that compares the incoming route with the prefix values for other aggregation rules.

For more information on BGP, refer to How to Configure BGP Tech Note.

The firewall provides a complete BGP implementation, which includes the following features:

- Specification of one BGP routing instance per virtual router.
- Routing policies based on route-map to control import, export and advertisement, prefix-based filtering, and address aggregation.
- Advanced BGP features that include route reflector, AS confederation, route flap dampening, and graceful restart.
- IGP-BGP interaction to inject routes to BGP using redistribution profiles.

BGP configuration consists of the following elements:

- Per-routing-instance settings, which include basic parameters such as local route ID and local AS and advanced options such as path selection, route reflector, AS confederation, route flap, and dampening profiles.
- Authentication profiles, which specify the MD5 authentication key for BGP connections.
- Peer group and neighbor settings, which include neighbor address and remote AS and advanced options such as neighbor attributes and connections.
- Routing policy, which specifies rule sets that peer groups and peers use to implement imports, exports, conditional advertisements, and address aggregation controls.

Perform the following procedure to configure BGP.

| **Configure BGP** | | |
|---|---|---|
| Step 1 | Configure general virtual router configuration settings. | See Virtual Routers for details. |

 © Palo Alto Networks, Inc.

| Configure BGP (Continued) | | |
|---|---|---|
| Step 2 | Configure standard BGP configuration settings. | 1. Select the **BGP** tab.<br>2. Select the **Enable** check box to enable the BGP protocol.<br>3. Assign an IP address to the virtual router in the **Router ID** box.<br>4. Enter the number of the AS to which the virtual router belongs in the **AS Number** box, based on the router ID. Range is 1-4294967295. |
| Step 3 | Configure general BGP configuration settings. | 1. Select **BGP > General**.<br>2. Select the **Reject Default Route** check box to ignore any default routes that are advertised by BGP peers.<br>3. Select the **Install Route** check box to install BGP routes in the global routing table.<br>4. Select the **Aggregate MED** check box to enable route aggregation even when routes have different Multi-Exit Discriminator (MED) values.<br>5. Enter a value for the **Default Local Preference** that specifies a value than can be used to determine preferences among different paths.<br>6. Select one of the following values for the AS format for interoperability purposes:<br>• 2 Byte (default value)<br>• 4 Byte<br>7. Enable or disable each of the following values for **Path Selection**:<br>• **Always Compare MED**—Enable this comparison to choose paths from neighbors in different autonomous systems.<br>• **Deterministic MED Comparison**—Enable this comparison to choose between routes that are advertised by IBGP peers (BGP peers in the same autonomous system).<br>8. Click **Add** to include a new authentication profile and configure the following settings:<br>• **Profile Name**—Enter a name to identify the profile.<br>• **Secret/Confirm Secret**—Enter and confirm a passphrase for BGP peer communications. |

| Configure BGP (Continued) | |
|---|---|
| Step 4 Configure BGP Advanced settings (Optional) | 1. On the **Advanced** tab, select **Graceful Restart** and configure the following timers:<br>• **Stale Route Time (sec)**—Specifies the length of time in seconds that a route can say in the stale state (range is 1- 3600; default is 120).<br>• **Local Restart Time (sec)**—Specifies the length of time in seconds that the local device waits to restart. This value is advertised to peers (range is 1-3600' default is 120).<br>• **Max Peer Restart Time (sec)**—Specifies the maximum length of time in seconds that the local device accepts as a grave period restart time for peer devices (range is 1-3600; default is 120).<br>2. Specify an IPv4 identifier to represent the reflector cluster in the **Reflector Cluster ID** box.<br>3. Specify the identifier for the AS confederation to be presented as a single AS to external BGP peers in the **Confederation Member AS** box.<br>4. Click **Add** and enter the following information for each Dampening Profile that you want to configure, select **Enable**, and click **OK**:<br>• **Profile Name**—Enter a name to identify the profile.<br>• **Cutoff**—Specify a route withdrawal threshold above which a route advertisement is suppressed (range is 0.0-1000.0; default: is 1.25).<br>• **Reuse**—Specify a route withdrawal threshold below which a suppressed route is used again (range is 0.0-1000.0; default is 5).<br>• **Max Hold Time (sec)**—Specify the maximum length of time in seconds that a route can be suppressed, regardless of how unstable it has been (range is 0-3600 seconds; default is 900).<br>• **Decay Half Life Reachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered reachable (range is 0-3600 seconds; default: is 300).<br>• **Decay Half Life Unreachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered unreachable (range is 0-3600; default is 300).<br>5. Click **OK**. |


© Palo Alto Networks, Inc.

| Configure BGP (Continued) | | |
|---|---|---|
| Step 5 | Configure the BGP peer group. | 1. Select the **Peer Group** tab and click **Add**.<br>2. Enter a **Name** for the peer group and select **Enable**.<br>3. Select the **Aggregated Confed AS Path** check box to include a path to the configured aggregated confederation AS.<br>4. Select the **Soft Reset with Stored Info** check box to perform a soft reset of the firewall after updating the peer settings.<br>5. Specify the type of peer or group from the **Type** drop-down and configure the associated settings (see below in this table for descriptions of Import Next Hop and Export Next Hop).<br>• **IBGP—Export Next Hop:** Specify **Original** or **Use self**<br>• **EBGP Confed—Export Next Hop:** Specify **Original** or **Use self**<br>• **EBGP Confed—Export Next Hop:** Specify **Original** or **Use self**<br>• **EBGP—Import Next Hop:** Specify **Original** or **Use self**, **Export Next Hop:** Specify **Resolve** or **Use self**. Select **Remove Private AS** if you want to force BGP to remove private AS numbers.<br>6. Click **OK** to save. |
| Step 6 | Configure Import and Export rules.<br>The import/export rules are used to import/export routes from/to other routers. For example, importing the default route from your Internet Service Provider. | 1. Select the **Import** tab and then click **Add** and enter a name in the **Rules** field and select the **Enable** check box.<br>2. Click **Add** and select the **Peer Group** to which the routes will be imported from.<br>3. Click the **Match** tab and define the options used to filter routing information. You can also define the Multi-Exit Discriminator (MED) value and a next hop value to routers or subnets for route filtering. The MED option is an external metric that lets neighbors know about the preferred path into an AS. A lower value is preferred over a higher value.<br>4. Click the **Action** tab and define the action that should occur (allow/deny) based on the filtering options defined in the **Match** tab. If **Deny** is selected, no further options need to be defined. If the **Allow** action is selected, define the other attributes.<br>5. Click the **Export** tab and define export attributes, which are similar to the **Import** settings, but are used to control route information that is exported from the firewall to neighbors.<br>6. Click **OK** to save. |

| Configure BGP (Continued) | | |
|---|---|---|
| Step 7 | Configure conditional advertising, which allows you to control what route to advertise in the event that a different route is not available in the local BGP routing table (LocRIB), indicating a peering or reachability failure. This is useful in cases where you want to try to force routes to one AS over another, for example if you have links to the Internet through multiple ISPs and you want traffic to be routed to one provider instead of the other unless there is a loss of connectivity to the preferred provider. | 1. Select the **Conditional Adv** tab, click **Add** and enter a name in the **Policy** field.<br>2. Select the **Enable** check box.<br>3. Click **Add** and in the **Used By** section enter the peer group(s) that will use the conditional advertisement policy.<br>4. Select the **Non Exist Filter** tab and define the network prefix(es) of the preferred route. This specifies the route that you want to advertise, if it is available in the local BGP routing table. If a prefix is going to be advertised and matches a Non Exist filter, the advertisement will be suppressed.<br>5. Select the **Advertise Filters** tab and define the prefix(es) of the route in the Local-RIB routing table that should be advertised in the event that the route in the non-exist filter is not available in the local routing table. If a prefix is going to be advertised and does not match a Non Exist filter, the advertisement will occur. |
| Step 8 | Configure aggregate options to summaries routes in the BGP configuration.<br><br>BGP route aggregation is used to control how BGP aggregates addresses. Each entry in the table results in one aggregate address being created. This will result in an aggregate entry in the routing table when at least one or more specific route matching the address specified is learned. | 1. Select the **Aggregate** tab, click **Add** and enter a name for the aggregate address.<br>2. In the **Prefix** field, enter the network prefix that will be the primary prefix for the aggregated prefixes.<br>3. Select the **Suppress Filters** tab and define the attributes that will cause the matched routes to be suppressed.<br>4. Select the **Advertise Filters** tab and define the attributes that will cause the matched routes to always be advertised to peers. |
| Step 9 | Configure redistribution rules.<br><br>This rule is used to redistribute host routes and unknown routes that are not on the local RIB to the peers routers. | 1. Select the **Redist Rules** tab and click **Add**.<br>2. In the **Name** field, enter an IP subnet or select a redistribution profile. You can also configure a new redistribution profile from the drop-down if needed.<br>3. Click the **Enable** check box to enable the rule.<br>4. In the **Metric** field, enter the route metric that will be used for the rule.<br>5. In the **Set Origin** drop-down, select **incomplete**, **igp**, or **egp**.<br>6. Optionally set MED, local preference, AS path limit and community values. |


© Palo Alto Networks, Inc.

# Session Settings and Timeouts

This section describes the global settings that affect TCP, UDP, and ICMPv6 sessions, in addition to IPv6, NAT64, NAT oversubscription, jumbo frame size, MTU, accelerated aging, and captive portal authentication. There is also a setting (Rematch Sessions) that allows you to apply newly configured security policies to sessions that are already in progress.

The **Device > Setup > Session** tab is where these session settings and timeouts are configured.

The first few topics below provide brief summaries of the Transport Layer of the OSI model, TCP, UDP, and ICMP. For more information about the protocols, refer to their respective RFCs. The remaining topics describe the session timeouts and settings.

- Transport Layer Sessions
- TCP
- UDP
- ICMP
- Configure Session Timeouts
- Configure Session Settings
- Prevent TCP Split Handshake Session Establishment

© Palo Alto Networks, Inc.

## Transport Layer Sessions

A network session is an exchange of messages that occurs between two or more communication devices, lasting for some period of time. A session is established and is torn down when the session ends. Different types of sessions occur at three layers of the OSI model: the Transport layer, the Session layer, and the Application layer.

The Transport Layer operates at Layer 4 of the OSI model, providing reliable or unreliable, end-to-end delivery and flow control of data. Internet protocols that implement sessions at the Transport layer include Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

 © Palo Alto Networks, Inc.

# TCP

Transmission Control Protocol (TCP) (RFC 793) is one of the main protocols in the Internet Protocol (IP) suite, and is so prevalent that it is frequently referenced together with IP as *TCP/IP*. TCP is considered a reliable transport protocol because it provides error-checking while transmitting and receiving segments, acknowledges segments received, and reorders segments that arrive in the wrong order. TCP also requests and provides retransmission of segments that were dropped. TCP is stateful and connection-oriented, meaning a connection between the sender and receiver is established for the duration of the session. TCP provides flow control of packets, so it can handle congestion over networks.

TCP performs a handshake during session setup to initiate and acknowledge a session. After the data is transferred, the session is closed in an orderly manner, where each side transmits a FIN packet and acknowledges it with an ACK packet. The handshake that initiates the TCP session is often a three-way handshake (an exchange of three messages) between the initiator and the listener, or it could be a variation, such as a four-way or five-way split handshake or a simultaneous open. The TCP Split Handshake Drop explains how to Prevent TCP Split Handshake Session Establishment.

Applications that use TCP as their transport protocol include Hypertext Transfer Protocol (HTTP), HTTP Secure (HTTPS), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), Telnet, Post Office Protocol version 3 (POP3), Internet Message Access Protocol (IMAP), and Secure Shell (SSH).

The following topics describe details of the PAN-OS implementation of TCP.

- TCP Half Closed and TCP Time Wait Timers
- Unverified RST Timer
- TCP Split Handshake Drop

## TCP Half Closed and TCP Time Wait Timers

The TCP connection termination procedure uses a TCP Half Closed timer, which is triggered by the first FIN the firewall sees for a session. The timer is named TCP Half Closed because only one side of the connection has sent a FIN. A second timer, TCP Time Wait, is triggered by the second FIN or a RST.

If the firewall were to have only one timer triggered by the first FIN, a setting that was too short could prematurely close the half-closed sessions. Conversely, a setting that was too long would make the session table grow too much and possibly use up all of the sessions. Two timers allow you to have a relatively long TCP Half Closed timer and a short TCP Time Wait timer, thereby quickly aging fully closed sessions and controlling the size of the session table.

The following figure illustrates when the firewall's two timers are triggered during the TCP connection termination procedure.

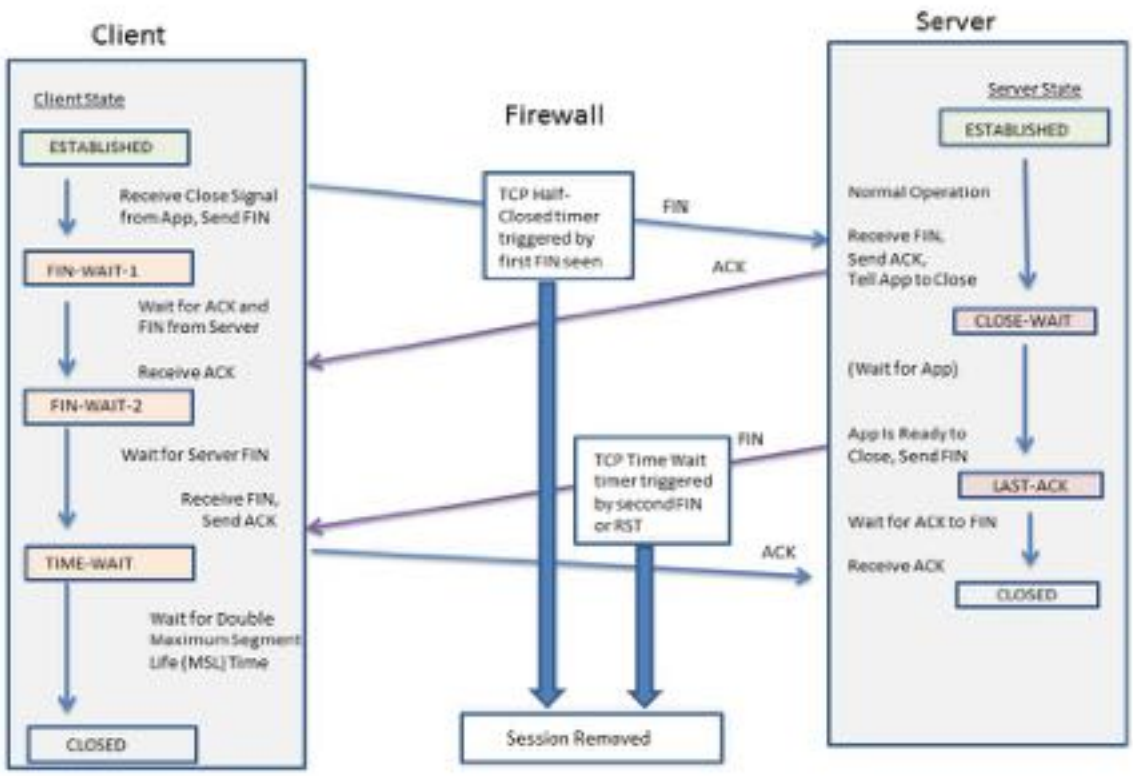


The TCP Time Wait timer should be set to a value less than the TCP Half Closed timer for the following reasons:

- The longer time allowed after the first FIN is seen gives the opposite side of the connection time to fully close the session.
- The shorter Time Wait time is because there is no need for the session to remain open for a long time after the second FIN or a RST is seen. A shorter Time Wait time frees up resources sooner, yet still allows time for the firewall to see the final ACK and possible retransmission of other datagrams.

If you configure a TCP Time Wait timer to a value greater than the TCP Half Closed timer, the commit will be accepted, but in practice the TCP Time Wait timer will not exceed the TCP Half Closed value.

The timers can be set globally or per application. The global settings are used for all applications by default. If you configure TCP wait timers at the application level, they override the global settings.

## Unverified RST Timer

If the firewall receives a Reset (RST) packet that cannot be verified (because it has an unexpected sequence number within the TCP window or it is from an asymmetric path), the Unverified RST timer controls the aging out of the session. It defaults to 30 seconds; the range is 1-600 seconds. The Unverified RST timer provides an additional security measure, explained in the second bullet below.

A RST packet will have one of three possible outcomes:

- A RST packet that falls outside the TCP window is dropped.

 © Palo Alto Networks, Inc.

- A RST packet that falls inside the TCP window but does not have the exact expected sequence number is unverified and subject to the Unverified RST timer setting. This behavior helps prevent denial of service (DoS) attacks where the attack tries to disrupt existing sessions by sending random RST packets to the firewall.
- A RST packet that falls within the TCP window and has the exact expected sequence number is subject to the TCP Time Wait timer setting.

## TCP Split Handshake Drop

The **Split Handshake** option in a Zone Protection profile will prevent a TCP session from being established if the session establishment procedure does not use the well-known three-way handshake, but instead uses a variation, such as a four-way or five-way split handshake or a simultaneous open.

The Palo Alto Networks next-generation firewall correctly handles sessions and all Layer 7 processes for split handshake and simultaneous open session establishment without enabling the **Split Handshake** option. Nevertheless, the **Split Handshake** option (which causes a TCP split handshake drop) is made available. When the **Split Handshake** option is configured for a Zone Protection profile and that profile is applied to a zone, TCP sessions for interfaces in that zone must be established using the standard three-way handshake; variations are not allowed.

The **Split Handshake** option is disabled by default.

The following illustrates the standard three-way handshake used to establish a TCP session with a PAN-OS firewall between the initiator (typically a client) and the listener (typically a server).

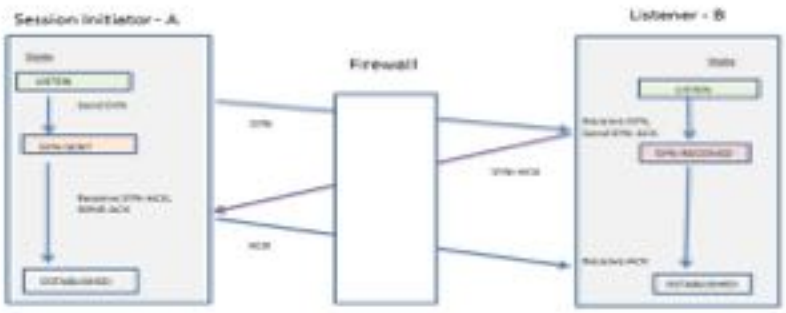


The **Split Handshake** option is configured for a Zone Protection profile that is assigned to a zone. An interface that is a member of the zone drops any synchronization (SYN) packets sent from the server, preventing the following variations of handshakes. The letter A in the figure indicates the session initiator and B indicates the listener. Each numbered segment of the handshake has an arrow indicating the direction of the segment from the sender to the receiver, and each segment indicates the control bit(s) setting.

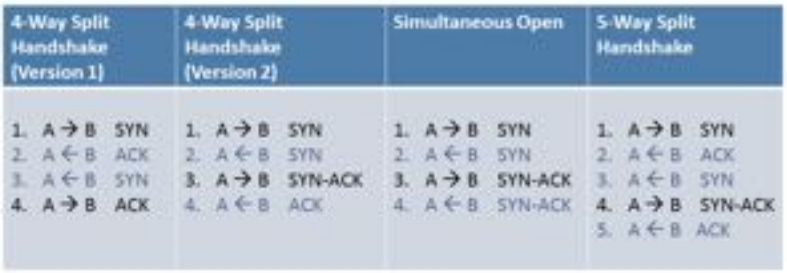

| 4-Way Split Handshake (Version 1) | 4-Way Split Handshake (Version 2) | Simultaneous Open | 5-Way Split Handshake |
|---|---|---|---|
| 1. A → B SYN | 1. A → B SYN | 1. A → B SYN | 1. A → B SYN |
| 2. A ← B ACK | 2. A ← B SYN | 2. A ← B SYN | 2. A ← B ACK |
| 3. A ← B SYN | 3. A → B SYN-ACK | 3. A → B SYN-ACK | 3. A ← B SYN |
| 4. A → B ACK | 4. A ← B ACK | 4. A ← B SYN-ACK | 4. A → B SYN-ACK |
| | | | 5. A ← B ACK |

You can Prevent TCP Split Handshake Session Establishment.

 © Palo Alto Networks, Inc.

## UDP

User Datagram Protocol (UDP) (RFC 768) is another main protocol of the IP suite, and is an alternative to TCP. UDP is stateless and connectionless in that there is no handshake to set up a session, and no connection between the sender and receiver; the packets may take different routes to get to a single destination. UDP is considered an unreliable protocol because it does not provide acknowledgments, error-checking, retransmission, or reordering of datagrams. Without the overhead required to provide those features, UDP has reduced latency and is faster than TCP. UDP is referred to as a best-effort protocol because there is no mechanism or guarantee to ensure that the data will arrive at its destination.

Although UDP uses a checksum for data integrity, it performs no error checking at the network interface level. Error checking is assumed to be unnecessary or is performed by the application rather than UDP itself. UDP has no mechanism to handle flow control of packets.

UDP is often used for applications that require faster speeds and time-sensitive, real-time delivery, such as Voice over IP (VoIP), streaming audio and video, and online games. UDP is transaction-oriented, so it is also used for applications that respond to small queries from many clients, such as Domain Name System (DNS) and Trivial File Transfer Protocol (TFTP).

© Palo Alto Networks, Inc.

# ICMP

Internet Control Message Protocol (ICMP) (RFC 792) is another one of the main protocols of the Internet Protocol suite; it operates at the Network layer of the OSI model. ICMP is used for diagnostic and control purposes, to send error messages about IP operations, or messages about requested services or the reachability of a host or router. Network utilities such as traceroute and ping are implemented by using various ICMP messages.

ICMP is a connectionless protocol that does not open or maintain actual sessions. However, the ICMP messages between two devices can be considered a session.

Palo Alto Networks firewalls support ICMPv4 and ICMPv6. ICMPv4 and ICMPv6 error packets can be controlled by configuring a security policy for a zone, and selecting the **icmp** or **ipv6-icmp** application in the policy. Additionally, the ICMPv6 error packet rate can be controlled through the session settings, as described in the section Configure Session Settings.

## ICMPv6 Rate Limiting

ICMPv6 rate limiting is a throttling mechanism to prevent flooding and DDoS attempts. The implementation employs an error packet rate and a token bucket, which work together to enable throttling and ensure that ICMP packets do not flood the network segments protected by the firewall.

First the global ICMPv6 error packet rate controls the rate at which ICMP error packets are allowed through the firewall; the default is 100 packets per second; the range is 10 to 65535 packets per second. If the firewall reaches the ICMP error packet rate, then the token bucket comes into play and throttling occurs, as follows.

The concept of a logical token bucket controls the rate at which ICMP messages can be transmitted. The number of tokens in the bucket is configurable, and each token represents an ICMP message that can be sent. The token count is decremented each time an ICMP message is sent; when the bucket reaches zero tokens, no more ICMP messages can be sent until another token is added to the bucket. The default size of the token bucket is 100 tokens (packets); the range is 10 to 65535 tokens.

To change the default token bucket size or error packet rate, see the section Configure Session Settings.

 © Palo Alto Networks, Inc.

## Configure Session Timeouts

A session timeout defines the duration of time for which PAN-OS maintains a session on the firewall after inactivity in the session. By default, when the session timeout for the protocol expires, PAN-OS closes the session.

On the firewall, you can define a number of timeouts for TCP, UDP, and ICMP sessions in particular. The Default timeout applies to any other type of session. All of these timeouts are global, meaning they apply to all of the sessions of that type on the firewall.

In addition to the global settings, you have the flexibility to define timeouts for an individual application in the **Objects > Applications** tab. The firewall applies application timeouts to an application that is in established state. When configured, timeouts for an application override the global TCP or UDP session timeouts.

Returning to the global settings, perform the optional tasks below if you need to change default values of the global session timeout settings for TCP, UDP, ICMP, Captive Portal authentication, or other types of sessions. All values are in seconds.

The defaults are optimal values. However, you can modify these according to your network needs. Setting a value too low could cause sensitivity to minor network delays and could result in a failure to establish connections with the firewall. Setting a value too high could delay failure detection.

| **Change Session Timeouts** | | |
|---|---|---|
| Step 1 | Access the Session Settings. | 1. Select **Device > Setup > Session**.<br>2. In the Session Timeouts section, click the Edit icon. |
| Step 2 | (Optional) Change miscellaneous timeouts. | • **Default**—Maximum length of time that a non-TCP/UDP or non-ICMP session can be open without a response (range is 1-1599999; default is 30).<br>• **Discard Default**—Maximum length of time that a non-TCP/UDP session remains open after PAN-OS denies a session based on security policies configured on the firewall (range is 1-1599999; default is 60).<br>• **Scan**—Maximum length of time that any session remains open after it is considered inactive; an application is regarded as inactive when it exceeds the application trickling threshold defined for the application (range is 5-30; default is 10).<br>• **Captive Portal**—Authentication session timeout for the Captive Portal web form. To access the requested content, the user must enter the authentication credentials in this form and be successfully authenticated (range is 1-1599999; default is 30).<br>To define other Captive Portal timeouts, such as the idle timer and the expiration time before the user must be re-authenticated, select **Device > User Identification > Captive Portal Settings**. See Configure Captive Portal in User-ID. |

| Change Session Timeouts (Continued) | |
|---|---|
| Step 3 (Optional) Change TCP timeouts. | • **Discard TCP**—Maximum length of time that a TCP session remains open after it is denied based on a security policy configured on the firewall. Default: 90. Range: 1-1599999.<br>• **TCP**—Maximum length of time that a TCP session remains open without a response, after a TCP session is in the Established state (after the handshake is complete and/or data is being transmitted). Default: 3600. Range: 1-1599999.<br>• **TCP Handshake**—Maximum length of time permitted between receiving the SYN-ACK and the subsequent ACK to fully establish the session. Default: 10. Range: 1-60.<br>• **TCP init**—Maximum length of time permitted between receiving the SYN and SYN-ACK prior to starting the TCP handshake timer. Default: 5. Range: 1-60.<br>• **TCP Half Closed**—Maximum length of time between receiving the first FIN and receiving the second FIN or a RST. Default: 120. Range: 1-604800.<br>• **TCP Time Wait**—Maximum length of time after receiving the second FIN or a RST. Default: 15. Range: 1-600.<br>• **Unverified RST**—Maximum length of time after receiving a RST that cannot be verified (the RST is within the TCP window but has an unexpected sequence number, or the RST is from an asymmetric path). Default: 30. Range: 1-600.<br>• See also the **Scan** timeout in the section (Optional) Change miscellaneous timeouts. |
| Step 4 (Optional) Change UDP timeouts. | • **Discard UDP**—Maximum length of time that a UDP session remains open after it is denied based on a security policy configured on the firewall. Default: 60. Range: 1-1599999.<br>• **UDP**—Maximum length of time that a UDP session remains open without a UDP response. Default: 30. Range: 1-1599999.<br>• See also the **Scan** timeout in the section (Optional) Change miscellaneous timeouts. |
| Step 5 (Optional) Change ICMP timeouts. | • **ICMP**—Maximum length of time that an ICMP session can be open without an ICMP response. Default: 6. Range: 1-1599999.<br>• See also the **Discard Default** and **Scan** timeout in the section (Optional) Change miscellaneous timeouts. |
| Step 6 Commit the changes. | Click **OK** and **Commit** the changes. |

 © Palo Alto Networks, Inc.

## Configure Session Settings

This topic describes various settings for sessions other than timeouts values. Perform these optional tasks if you need to change the default settings.

| Change Session Settings | | |
|---|---|---|
| Step 1 | Access the Session Settings. | 1. Select **Device > Setup > Session.**<br>2. In the **Session Settings** section, click the **Edit** icon. |
| Step 2 | (Optional) Change session settings. | • **Rematch Sessions**—Causes the firewall to apply newly configured security policies to sessions that are already in progress. This capability is enabled by default. If this setting is disabled, any policy change applies only to sessions initiated after the policy change was committed.<br>For example, if a Telnet session started while an associated policy was configured that allowed Telnet, and you subsequently committed a policy change to deny Telnet, the firewall applies the revised policy to the current session and blocks it.<br>• **ICMPv6 Token Bucket Size**—Default: 100 tokens. See the section ICMPv6 Rate Limiting.<br>• **ICMPv6 Error Packet Rate (per sec)**—Default: 100. See the section ICMPv6 Rate Limiting.<br>• **Enable IPv6 Firewalling**—Enables firewall capabilities for IPv6. All IPv6-based configurations are ignored if IPv6 is not enabled. Even if IPv6 is enabled for an interface, the **IPv6 Firewalling** setting must also be enabled for IPv6 to function.<br>• **Enable Jumbo Frame**—Enables jumbo frame support on Ethernet interfaces. Jumbo frames have a maximum transmission unit (MTU) of 9216 bytes and are available on certain platforms.<br>**Global MTU**—Establishes the maximum transmission unit (MTU) available globally; this field is tied to the **Enable Jumbo Frame** field.<br>  • If you do not check **Enable Jumbo Frame**, the **Global MTU** defaults to 1500 bytes; the range is 576 to 1500 bytes.<br>  • If you check **Enable Jumbo Frame**, the **Global MTU** defaults to 9192 bytes; the range is 9192 to 9216 bytes.<br> If you enable jumbo frames and you have interfaces where the MTU is not specifically configured, those interfaces will automatically inherit the jumbo frame size. Therefore, before you enable jumbo frames, if you have any interface that you do not want to have jumbo frames, you must set the MTU for that interface to 1500 bytes or another value.<br>• **NAT64 IPv6 Minimum Network MTU**—Sets the global MTU for IPv6 translated traffic. The default of 1280 bytes is based on the standard minimum MTU for IPv6 traffic. |

| Change Session Settings (Continued) | |
|---|---|
| | • **NAT Oversubscription Rate**—If NAT is configured to be Dynamic IP and Port (DIPP) translation, an oversubscription rate can be configured to multiply the number of times that the same translated IP address and port pair can be used concurrently. The rate is 1, 2, 4, or 8. The default setting is based on the firewall platform.<br>  • A rate of 1 means no oversubscription; each translated IP address and port pair can be used only once at a time.<br>  • If the setting is **Platform Default**, user configuration of the rate is disabled and the default oversubscription rate for the platform applies.<br>Reducing the oversubscription rate decreases the number of source device translations, but provides higher NAT rule capacities.<br>• **Accelerated Aging**—Enables faster aging-out of idle sessions. Select the check box to enable accelerated aging, and, if necessary, change the threshold (%) and scaling factor.<br>  • **Accelerated Aging Threshold**—Percentage of the session table that is full when accelerated aging begins. The default is 80%. When the session table reaches this threshold (% full), PAN-OS applies the **Accelerated Aging Scaling Factor** to the aging calculations for all sessions.<br>  • **Accelerated Aging Scaling Factor**—Scaling factor used in the accelerated aging calculations. The default scaling factor is 2, meaning that the accelerated aging occurs at a rate twice as fast as the configured idle time. The configured idle time divided by 2 results in a faster timeout of one-half the time. To calculate the session's accelerated aging, PAN-OS divides the configured idle time (for that type of session) by the scaling factor to determine a shorter timeout.<br>For example, if the scaling factor is 10, a session that would normally time out after 3600 seconds would time out 10 times faster (in 1/10 of the time), which is 360 seconds. |
| Step 3 Commit the changes. | Click **OK** and **Commit** the changes. |

 © Palo Alto Networks, Inc.

## Prevent TCP Split Handshake Session Establishment

You can configure a TCP Split Handshake Drop in a Zone Protection profile to prevent TCP sessions from being established unless they use the standard three-way handshake. This task assumes that the interface where this feature is desired is already assigned to a security zone.

| Configure a Zone Protection Profile to Prevent TCP Split Handshake Sessions | | |
|---|---|---|
| Step 1 | Configure a Zone Protection profile to prevent TCP sessions that use anything other than a three-way handshake to establish a session. | 1. Select **Network > Network Profiles > Zone Protection** and click **Add** to create a new profile (or select an existing profile).<br>2. If creating a new profile, enter a **Name** for the profile and an optional **Description**.<br>3. Select **Packet Based Attack Protection > TCP Drop** and select the **Split Handshake** check box.<br>4. Click **OK**. |
| Step 2 | Apply the profile to one or more security zones. | 1. Select **Network > Zones** and select the zone where you want to assign the zone protection profile.<br>2. In the Zone window, from the **Zone Protection Profile** drop-down, select the profile you configured in Step 1.<br>Alternatively, you could start creating a new profile here by clicking **Zone Protection Profile**, in which case you would continue accordingly.<br>3. Click **OK**.<br>4. (Optional) Repeat steps 1-3 to apply the profile to additional zones. |
| Step 3 | Save the configuration. | Click **OK** and **Commit**. |

# DHCP

This section describes Dynamic Host Configuration Protocol (DHCP) and the tasks required to configure an interface on a Palo Alto Networks firewall to act as a DHCP server, client, or relay agent. By assigning these roles to different interfaces, the firewall can perform multiple roles.


- DHCP Overview
- DHCP Messages
- DHCP Addressing
- DHCP Options
- Firewall as a DHCP Server and Client
- Configure an Interface as a DHCP Server
- Configure an Interface as a DHCP Client
- Configure an Interface as a DHCP Relay Agent
- Monitor and Troubleshoot DHCP


 © Palo Alto Networks, Inc.

## DHCP Overview

DHCP is a standardized protocol defined in RFC 2131, Dynamic Host Configuration Protocol. DHCP has two main purposes: to provide TCP/IP and link-layer configuration parameters and to provide network addresses to dynamically configured hosts on a TCP/IP network.

DHCP uses a client-server model of communication. This model consists of three roles that the device can fulfill: DHCP client, DHCP server, and DHCP relay agent.

- A device acting as a DHCP client (host) can request an IP address and other configuration settings from a DHCP server. Users on client devices save configuration time and effort, and need not know the network's addressing plan or other resources and options they are inheriting from the DHCP server.
- A device acting as a DHCP server can service clients. By using any of three DHCP Addressing mechanisms, the network administrator saves configuration time and has the benefit of reusing a limited number of IP addresses when a client no longer needs network connectivity. The server can deliver IP addressing and many DHCP options to many clients.
- A device acting as a DHCP relay agent transmits DHCP messages between DHCP clients and servers.

DHCP uses User Datagram Protocol (UDP), RFC 768, as its transport protocol. DHCP messages that a client sends to a server are sent to well-known port 67 (UDP—Bootstrap Protocol and DHCP). DHCP Messages that a server sends to a client are sent to port 68.

An interface on a Palo Alto Networks firewall can perform the role of a DHCP server, client, or relay agent. The interface of a DHCP server or relay agent must be a Layer 3 Ethernet, Aggregated Ethernet, or Layer 3 VLAN interface. You configure the firewall's interfaces with the appropriate settings for any combination of roles. The behavior of each role is summarized in Firewall as a DHCP Server and Client.

The firewall supports DHCPv4 Server and DHCPv6 Relay. However, a single interface cannot support both DHCPv4 Server and DHCPv6 Relay.

The DHCP standard, RFC 2131, is designed to support IPv4 and IPv6 addresses. The Palo Alto Networks implementations of DHCP server and DHCP client support IPv4 addresses only. Its DHCP relay implementation supports IPv4 and IPv6.

DHCP client is not supported in High Availability active/active mode.

## DHCP Messages

DHCP uses eight standard message types, which are identified by an option type number in the DHCP message. For example, when a client wants to find a DHCP server, it broadcasts a DHCPDISCOVER message on its local physical subnetwork. If there is no DHCP server on its subnet and if DHCP Helper or DHCP Relay is configured properly, the message is forwarded to DHCP servers on a different physical subnet. Otherwise, the message will go no further than the subnet on which it originated. One or more DHCP servers will respond with a DHCPOFFER message that contains an available network address and other configuration parameters.

When the client needs an IP address, it sends a DHCPREQUEST to one or more servers. Of course if the client is requesting an IP address, it doesn't have one yet, so RFC 2131 requires that the broadcast message the client sends out have a source address of 0 in its IP header.

When a client requests configuration parameters from a server, it might receive responses from more than one server. Once a client has received its IP address, it is said that the client has at least an IP address and possibly other configuration parameters *bound* to it. DHCP servers manage such binding of configuration parameters to clients.

The following table lists the DHCP messages.

| DHCP Message | Description |
|---|---|
| DHCPDISCOVER | Client broadcast to find available DHCP servers. |
| DHCPOFFER | Server response to client's DHCPDISCOVER, offering configuration parameters. |
| DHCPREQUEST | Client message to one or more servers to do any of the following:<br>• Request parameters from one server and implicitly decline offers from other servers.<br>• Confirm that a previously allocated address is correct after, for example, a system reboot.<br>• Extend the lease of a network address. |
| DHCPACK | Server to client acknowledgment message containing configuration parameters, including a confirmed network address. |
| DHCPNAK | Server to client negative acknowledgment indicating the client's understanding of the network address is incorrect (for example, if the client has moved to a new subnet), or a client's lease has expired. |
| DHCPDECLINE | Client to server message indicating the network address is already being used. |
| DHCPRELEASE | Client to server message giving up the user of the network address and canceling the remaining time on the lease. |
| DHCPINFORM | Client to server message requesting only local configuration parameters; client has an externally configured network address. |

 © Palo Alto Networks, Inc.

# DHCP Addressing

▲ DHCP Address Allocation Methods
▲ DHCP Leases

## DHCP Address Allocation Methods

There are three ways that a DHCP server either assigns or sends an IP address to a client:

- **Automatic allocation**—The DHCP server assigns a permanent IP address to a client from its **IP Pools**. On the firewall, a **Lease** specified as **Unlimited** means the allocation is permanent.
- **Dynamic allocation**—The DHCP server assigns a reusable IP address from **IP Pools** of addresses to a client for a maximum period of time, known as a *lease*. This method of address allocation is useful when the customer has a limited number of IP addresses; they can be assigned to clients who need only temporary access to the network. See the DHCP Leases section.
- **Static allocation**—The network administrator chooses the IP address to assign to the client and the DHCP server sends it to the client. A static DHCP allocation is permanent; it is done by configuring a DHCP server and choosing a **Reserved Address** to correspond to the **MAC Address** of the client device. The DHCP assignment remains in place even if the client logs off, reboots, has a power outage, etc.

  Static allocation of an IP address is useful, for example, if you have a printer on a LAN and you do not want its IP address to keep changing, because it is associated with a printer name through DNS. Another example is if a client device is used for something crucial and must keep the same IP address, even if the device is turned off, unplugged, rebooted, or a power outage occurs, etc.

  Keep these points in mind when configuring a **Reserved Address**:

  - It is an address from the **IP Pools**. You may configure multiple reserved addresses.
  - If you configure no **Reserved Address**, the clients of the server will receive new DHCP assignments from the pool when their leases expire or if they reboot, etc. (unless you specified that a **Lease** is **Unlimited**).
  - If you allocate all of the addresses in the **IP Pools** as a **Reserved Address,** there are no dynamic addresses free to assign to the next DHCP client requesting an address.
  - You may configure a **Reserved Address** without configuring a **MAC Address**. In this case, the DHCP server will not assign the **Reserved Address** to any device. You might reserve a few addresses from the pool and statically assign them to a fax and printer, for example, without using DHCP.

## DHCP Leases

A lease is defined as the time period for which a DHCP server allocates a network address to a client. The lease might be extended (renewed) upon subsequent requests. If the client no longer needs the address, it can release the address back to the server before the lease is up. The server is then free to assign that address to a different client if it has run out of unassigned addresses.

© Palo Alto Networks, Inc.

The lease period configured for a DHCP server applies to all of the addresses that a single DHCP server (interface) dynamically assigns to its clients. That is, all of that interface's addresses assigned dynamically are of **Unlimited** duration or have the same **Timeout** value. A different DHCP server configured on the firewall may have a different lease term for its clients. A **Reserved Address** is a static address allocation and is not subject to the lease terms.

Per the DHCP standard, RFC 2131, a DHCP client does not wait for its lease to expire, because it risks getting a new address assigned to it. Instead, when a DHCP client reaches the halfway point of its lease period, it attempts to extend its lease so that it retains the same IP address. Thus, the lease duration is like a sliding window.

Typically if an IP address was assigned to a device, the device was subsequently taken off the network and its lease was not extended, the DHCP server will let that lease run out. Because the client is gone from the network and no longer needs the address, the lease duration in the server is reached and the lease is in "Expired" state.

The firewall has a hold timer that prevents the expired IP address from being reassigned immediately. This behavior temporarily reserves the address for the device in case it comes back onto the network. But if the address pool runs out of addresses, the server re-allocates this expired address before the hold timer expires. Expired addresses are cleared automatically as the systems needs more addresses or when the hold timer releases them.

In the CLI, use the `show dhcp server lease` operational command to view lease information about the allocated IP addresses. If you do not want to wait for expired leases to be released automatically, you can use the `clear dhcp lease interface` *`value`* `expired-only` command to clear expired leases, making those addresses available in the pool again. You can use the `clear dhcp lease interface` *`value`* `ip` *`ip`* command to release a particular IP address. Use the `clear dhcp lease interface` *`value`* `mac` *`mac_address`* command to release a particular MAC address.

 © Palo Alto Networks, Inc.

# DHCP Options

The history of DHCP and DHCP options traces back to the Bootstrap Protocol (BOOTP). BOOTP was used by a host to configure itself dynamically during its booting procedure. A host could receive an IP address and a file from which to download a boot program from a server, along with the server's address and the address of an Internet gateway.

Included in the BOOTP packet was a vendor information field, which could contain a number of tagged fields containing various types of information, such as the subnet mask, the BOOTP file size, and many other values. RFC 1497 describes the BOOTP Vendor Information Extensions. DHCP replaces BOOTP; BOOTP is not supported on the firewall.

These extensions eventually expanded with the use of DHCP and DHCP host configuration parameters, also known as options. Similar to vendor extensions, DHCP options are tagged data items that provide information to a DHCP client. The options are sent in a variable-length field at the end of a DHCP message. For example, the DHCP Message Type is option 53, and a value of 1 indicates the DHCPDISCOVER message. DHCP options are defined in RFC 2132, DHCP Options and BOOTP Vendor Extensions.

A DHCP client can negotiate with the server, limiting the server to send only those options that the client requests.

- Predefined DHCP Options
- Multiple Values for a DHCP Option
- DHCP Options 43, 55, and 60 and Other Customized Options

## Predefined DHCP Options

Palo Alto Networks firewalls support user-defined and predefined DHCP options in the DHCP server implementation. Such options are configured on the DHCP server and sent to the clients that sent a DHCPREQUEST to the server. The clients are said to *inherit* and implement the options that they are programmed to accept.

The firewall supports the following predefined options on its DHCP servers, shown in the order in which they appear on the **DHCP Server** configuration screen:

| DHCP Option | DHCP Option Name |
|---|---|
| 51 | Lease duration |
| 3 | Gateway |
| 1 | IP Pool Subnet (mask) |
| 6 | Domain Name System (DNS) server address (primary and secondary) |
| 44 | Windows Internet Name Service (WINS) server address (primary and secondary) |
| 41 | Network Information Service (NIS) server address (primary and secondary) |
| 42 | Network Time Protocol (NTP) server address (primary and secondary) |
| 70 | Post Office Protocol Version 3 (POP3) server address |

| DHCP Option | DHCP Option Name |
|---|---|
| 69 | Simple Mail Transfer Protocol (SMTP) server address |
| 15 | DNS suffix |

As mentioned, you can also configure vendor-specific and customized options, which support a wide variety of office equipment, such as IP phones and wireless infrastructure devices. Each option code supports multiple values, which can be IP address, ASCII, or hexadecimal format. With the firewall enhanced DCHP option support, branch offices do not need to purchase and manage their own DHCP servers in order to provide vendor-specific and customized options to DHCP clients.

## Multiple Values for a DHCP Option

You can enter multiple option values for an **Option Code** with the same **Option Name**, but all values for a particular code and name combination must be the same type (IP address, ASCII, or hexadecimal). If one type is inherited or entered, and later a different type is entered for the same code and name combination, the second type will overwrite the first type.

You can enter an **Option Code** more than once by using a different **Option Name**. In this case, the **Option Type** for the Option Code can differ among the multiple option names. For example, if option Coastal Server (option code 6) is configured with IP address type, option Server XYZ (option code 6) with ASCII type is also allowed.

The firewall sends multiple values for an option (strung together) to a client in order from top to bottom. Therefore, when entering multiple values for an option, enter the values in the order of preference, or else move the options to achieve your preferred order in the list. The order of options in the firewall configuration determines the order that the options appear in DHCPOFFER and DHCPACK messages.

You can enter an option code that already exists as a predefined option code, and the customized option code will override the predefined DHCP option; the firewall issues a warning.

## DHCP Options 43, 55, and 60 and Other Customized Options

The following table describes the option behavior for several options described in RFC 2132.

| Option Code | Option Name | Option Description/Behavior |
|---|---|---|
| 43 | Vendor Specific Information | Sent from server to client. Vendor-specific information that the DHCP server has been configured to offer to the client. The information is sent to the client only if the server has a Vendor Class Identifier (VCI) in its table that matches the VCI in the client's DHCPREQUEST.<br>An Option 43 packet can contain multiple vendor-specific pieces of information. It can also include encapsulated, vendor-specific extensions of data. |
| 55 | Parameter Request List | Sent from client to server. List of configuration parameters (option codes) that a DHCP client is requesting, possibly in order of the client's preference. The server tries to respond with options in the same order. |

 © Palo Alto Networks, Inc.

| Option Code | Option Name | Option Description/Behavior |
|---|---|---|
| 60 | Vendor Class Identifier (VCI) | Sent from client to server. Vendor type and configuration of a DHCP client. The DHCP client sends option code 60 in a DHCPREQUEST to the DHCP server. When the server receives option 60, it sees the VCI, finds the matching VCI in its own table, and then it returns option 43 with the value (that corresponds to the VCI), thereby relaying vendor-specific information to the correct client. Both the client and server have knowledge of the VCI. |

You can send custom, vendor-specific option codes that are not defined in RFC 2132. The option codes can be in the range 1-254 and of fixed or variable length.

Custom DHCP options are not validated by the DHCP Server; you must ensure that you enter correct values for the options you create.

For ASCII and hexadecimal DHCP option types, the option value can be a maximum of 255 octets.

## Firewall as a DHCP Server and Client

The firewall can function as a DHCP server and as a DHCP client. The firewall DHCP server operates in the following manner:

- When the DHCP server receives a DHCPDISCOVER message from a client, the server replies with a DHCPOFFER message containing all of the predefined and user-defined options in the order they appear in the configuration. The client selects the options it needs and responds with a DHCPREQUEST message.
- When the server receives a DHCPREQUEST message from a client, the server replies with its DHCPACK message containing only the options specified in the request.
- Dynamic Host Configuration Protocol, RFC 2131, is designed to support IPv4 and IPv6 addresses. The Palo Alto Networks implementation of DHCP server supports IPv4 addresses only.

The firewall DHCP Client operates in the following manner:

- When the DHCP client receives a DHCPOFFER from the server, the client automatically caches all of the options offered for future use, regardless of which options it had sent in its DHCPREQUEST.
- By default and to save memory consumption, the client caches only the first value of each option code if it receives multiple values for a code.
- There is no maximum length for DHCP messages unless the DHCP client specifies a maximum in option 57 in its DHCPDISCOVER or DHCPREQUEST messages.

 © Palo Alto Networks, Inc.

## Configure an Interface as a DHCP Server

The prerequisites for this task are:

- Configure a Layer 3 Ethernet or Layer 3 VLAN interface.
- Assign the interface to a virtual router and a zone.
- Determine a valid pool of IP addresses from your network plan that you can designate to be assigned by your DHCP server to clients.
- Collect the DHCP options, values, and Vendor Class Identifiers you plan to configure.

Perform the following task to configure an interface on the firewall to act as a DHCP server. You can configure multiple DHCP servers.

| **Configure an Interface as a DHCP Server** | |
|---|---|
| Step 1 Select an interface to be a DHCP Server. | 1. Select **Network > DHCP > DHCP Server** and click **Add**.<br>2. Enter an **Interface** name or select one from the drop-down.<br>3. For **Mode**, select **enabled** or **auto** mode. Auto mode enables the server and disables it if another DHCP server is detected on the network. The **disabled** setting disables the server.<br>4. (Optional) Click the **Ping IP when allocating new IP** check box if you want the server to ping the IP address before it assigns that address to its client.<br>If the ping receives a response, that means a different device already has that address, so it is not available. The server assigns the next address from the pool instead. This behavior is similar to Optimistic Duplicate Address Detection (DAD) for IPv6, RFC 4429.<br>After you set options and return to the DHCP server tab, the **Probe IP** column for the interface indicates if this check box was checked. |

| Configure an Interface as a DHCP Server (Continued) | |
|---|---|
| Step 2 Configure the predefined DHCP Options that the server sends to its clients. | • In the Options section, select a **Lease** type:<br>  • **Unlimited** causes the server to dynamically choose IP addresses from the **IP Pools** and assign them permanently to clients.<br>  • **Timeout** determines how long the lease will last. Enter the number of **Days** and **Hours**, and optionally the number of **Minutes**.<br>• **Inheritance Source**—Leave **None** or select a source DHCP client interface or PPPoE client interface to propagate various server settings into the DHCP server. If you specify an **Inheritance Source**, select one or more options below that you want **inherited** from this source.<br>Specifying an inheritance source allows the firewall to quickly add DHCP options from the upstream server received by the DHCP client. It also keeps the client options updated if the source changes an option. For example, if the source replaces its NTP server (which had been identified as the **Primary NTP** server), the client will automatically inherit the new address as its **Primary NTP** server.<br> When inheriting DHCP option(s) that contain multiple IP addresses, the firewall uses only the first IP address contained in the option to conserve cache memory. If you require multiple IP addresses for a single option, configure the DHCP options directly on that firewall rather than configure inheritance.<br>• **Check inheritance source status**—If you selected an **Inheritance Source**, clicking this link opens the **Dynamic IP Interface Status** window, which displays the options that were inherited from the DHCP client.<br>• **Gateway**—IP address of the network gateway (an interface on the firewall) that is used to reach any device not on the same LAN as this DHCP server.<br>• **Subnet Mask**—Network mask used with the addresses in the **IP Pools**. |


© Palo Alto Networks, Inc.

| Configure an Interface as a DHCP Server (Continued) | |
|---|---|
| | For the following fields, click the down arrow and select **None**, or **inherited**, or enter a remote server's IP address that your DHCP server will send to clients for accessing that service. If you select **inherited,** the DHCP server inherits the values from the source DHCP client specified as the **Inheritance Source**.<br>• **Primary DNS, Secondary DNS**—IP address of the preferred and alternate Domain Name System (DNS) servers.<br>• **Primary WINS, Secondary WINS**—IP address of the preferred and alternate Windows Internet Naming Service (WINS) servers.<br>• **Primary NIS, Secondary NIS**—IP address of the preferred and alternate Network Information Service (NIS) servers.<br>• **Primary NTP, Secondary NTP**—IP address of the available Network Time Protocol servers.<br>• **POP3 Server**—IP address of a Post Office Protocol (POP3) server.<br>• **SMTP Server**—IP address of a Simple Mail Transfer Protocol (SMTP) server.<br>• **DNS Suffix**—Suffix for the client to use locally when an unqualified hostname is entered that it cannot resolve. |

| Configure an Interface as a DHCP Server (Continued) | |
|---|---|
| Step 3 (Optional) Configure a vendor-specific or custom DHCP option that the DHCP server sends to its clients. | 1. In the Custom DHCP Options section, click **Add** and enter a descriptive **Name** to identify the DHCP option.<br>2. Enter the **Option Code** you want to configure the server to offer (range is 1-254). (See RFC 2132 for option codes.)<br>3. If the **Option Code** is **43**, the **Vendor Class Identifier** field appears. Enter a VCI, which is a string or hexadecimal value (with 0x prefix) used as a match against a value that comes from the client Request containing option 60. The server looks up the incoming VCI in its table, finds it, and returns Option 43 and the corresponding option value.<br>4. **Inherit from DHCP server inheritance source**—Select this check box only if you specified an **Inheritance Source** for the DHCP Server predefined options and you want the vendor-specific and custom options also to be **inherited** from this source.<br>5. **Check inheritance source status**—If you selected an **Inheritance Source**, clicking this link opens **Dynamic IP Interface Status**, which displays the options that were inherited from the DHCP client.<br>6. If you did not select the **Inherit from DHCP server inheritance source** check box, select an **Option Type**: **IP Address**, **ASCII**, or **Hexadecimal**. Hexadecimal values must start with the 0x prefix.<br>7. Enter the **Option Value** you want the DHCP server to offer for that **Option Code**. You can enter multiple values on separate lines.<br>8. Click **OK.** |
| Step 4 (Optional) Add another vendor-specific or custom DHCP option. | 1. Repeat Step 3 to enter another custom DHCP Option.<br>• You can enter multiple option values for an **Option Code** with the same **Option Name**, but all values for an **Option Code** must be the same type (**IP Address**, **ASCII**, or **Hexadecimal**). If one type is inherited or entered and a different type is entered for the same **Option Code** and the same **Option Name**, the second type will overwrite the first type.<br>When entering multiple values for an option, enter the values in the order of preference, or else move the Custom DHCP Options to achieve the preferred order in the list. Select an option and click **Move Up** or **Move Down**.<br>• You can enter an **Option Code** more than once by using a different **Option Name**. In this case, the **Option Type** for the Option Code can differ among the multiple option names.<br>2. Click **OK.** |

 © Palo Alto Networks, Inc.

| Configure an Interface as a DHCP Server (Continued) | | |
|---|---|---|
| Step 5 | Identify the stateful pool of IP addresses from which the DHCP server chooses an address and assigns it to a DHCP client.<br> If you are not the network administrator for your network, ask the network administrator for a valid pool of IP addresses from the network plan that can be designated to be assigned by your DHCP server. | 1. In the **IP Pools** field, click **Add** and enter the range of IP addresses from which this server assigns an address to a client. Enter an IP subnet and subnet mask (for example, 192.168.1.0/24) or a range of IP addresses (for example, 192.168.1.10-192.168.1.20). At least one IP pool is required.<br>2. Optionally repeat to specify another IP address pool. |
| Step 6 | (Optional) Specify an IP address from the IP pools that will not be assigned dynamically. If you also specify a **MAC Address**, the **Reserved Address** is assigned to that device when the device requests an IP address through DHCP.<br> See the DHCP Addressing section for an explanation of allocation of a **Reserved Address**. | 1. In the **Reserved Address** field, click **Add.**<br>2. Enter an IP address from the **IP Pools** (format *x.x.x.x*) that you do not want to be assigned dynamically by the DHCP server.<br>3. (Optional) Specify the **MAC Address** (format *xx:xx:xx:xx:xx:xx*) of the device to which you want to permanently assign the IP address specified in Step 2.<br>4. (Optional) Repeat Steps 2-3 to reserve another address. |
| Step 7 | Save the configuration. | Click **OK** and **Commit** the change. |

## Configure an Interface as a DHCP Client

Before configuring a firewall interface as a DHCP Client, make sure you have configured a Layer 3 Ethernet or Layer 3 VLAN interface, and the interface is assigned to a virtual router and a zone. Perform this task if you need to use DHCP to request an IPv4 address for an interface on your firewall.

| **Configure an Interface as a DHCP Client** | |
|---|---|
| Step 1 Configure an interface as a DHCP client. | 1. Select **Network** > **Interfaces**.<br>2. On the **Ethernet** tab or the **VLAN** tab, click **Add** and enter an interface, or click a configured interface, that you want to be a DHCP client.<br>The following shows the **Ethernet Interface** screen:<br> <br>3. Click the **IPv4** tab; for **Type**, select **DHCP Client**.<br>4. Select the **Enable** check box.<br>5. Optionally select the **Automatically create default route pointing to default gateway provided by server** check box. This causes the firewall to create a static route to a default gateway that will be useful when clients are trying to access many destinations that do not need to have routes maintained in a routing table on the firewall.<br>6. Optionally enter a **Default Route Metric** (priority level) for the route between the firewall and the DHCP server (range is 1-65535; there is no default metric). A route with a lower number has higher priority during route selection. For example, a route with a metric of 10 is used before a route with a metric of 100.<br>7. Optionally click **Show DHCP Client Runtime Info** to see all of the settings the client has inherited from its DHCP server. |
| Step 2 Save the configuration. | Click **OK** and **Commit** the change.<br>Now the Ethernet interface indicates **Dynamic-DHCP Client** in its **IP Address** field on the **Ethernet** tab. |

 © Palo Alto Networks, Inc.

| Configure an Interface as a DHCP Client | | |
|---|---|---|
| Step 3 | (Optional) See which interfaces on the firewall are configured as DHCP clients. | 1. Select **Network > Interfaces > Ethernet** and look in the **IP Address** field to see which interfaces indicate DHCP Client.<br>2. Select **Network > Interfaces > VLAN** and look in the **IP Address** field to see which interfaces indicate DHCP Client. |

## Configure an Interface as a DHCP Relay Agent

Before configuring a firewall interface as a DHCP relay agent, make sure you have configured a Layer 3 Ethernet or Layer 3 VLAN interface, and the interface is assigned to a virtual router and a zone. You want that interface be able to pass DHCP messages between clients and servers. The interface can forward messages to a maximum of four external DHCP servers. A client DHCPDISCOVER message is sent to all configured servers, and the DHCPOFFER message of the first server that responds is relayed back to the requesting client.

Perform the following task to configure an interface on the firewall to act as a DHCP relay agent:

| **Configure an Interface as a DHCP Relay Agent** | | |
|---|---|---|
| Step 1 | Select DHCP Relay. | Select **Network** > **DHCP > DHCP Relay**. |
| Step 2 | Specify the IP address of each DHCP server with which the DHCP relay agent will communicate. | 1. In the **Interface** field, select from the drop-down the interface you want to be the DHCP relay agent.<br>2. Check either the **IPv4** or **IPv6** check box, indicating the type of DHCP server address you will specify.<br>3. If you checked **IPv4**, in the **DHCP Server IP Address** field, click **Add**. Enter the address of the DHCP server to and from which you will relay DHCP messages.<br>4. If you checked **IPv6**, in the **DHCP Server IPv6 Address** field, click **Add**. Enter the address of the DHCP server to and from which you will relay DHCP messages. If you specify a *multicast* address, also specify an outgoing **Interface**.<br>5. Optionally repeat Steps 2-4 to enter a maximum of four DHCP server addresses. |
| Step 3 | Save the configuration. | Click **OK** and **Commit** the change. |

 © Palo Alto Networks, Inc.

# Monitor and Troubleshoot DHCP

You can view the status of dynamic address leases that your DHCP server has assigned or that your DHCP client has been assigned by issuing commands from the CLI. You can also clear leases before they time out and are released automatically.

- View DHCP Server Information
- Clear Leases Before They Expire Automatically
- View DHCP Client Information
- Gather Debug Output about DHCP

## View DHCP Server Information

To view DHCP pool statistics, IP addresses the server has assigned, the corresponding MAC address, state and duration of the lease, and time the lease began, use the following command. If the address was configured as a **Reserved Address**, the `state` column indicates `reserved` and there is no `duration` or `lease_time`. If the lease was configured as **Unlimited**, the `duration` column displays a value of `0`.

```
admin@PA-200> show dhcp server lease all
interface: "ethernet1/2"
Allocated IPs: 1, Total number of IPs in pool: 5. 20.0000% used
ip              mac                state      duration    lease_time
192.168.3.11    f0:2f:af:42:70:cf  committed  0           Wed Jul  2 08:10:56 2014
admin@PA-200>
```

To view the options that a DHCP server has assigned to clients, use the following command:

```
admin@PA-200> show dhcp server settings all
Interface       GW            DNS1           DNS2         DNS-Suffix    Inherit source
------------------------------------------------------------------------------------------
ethernet1/2     192.168.3.1   10.43.2.10     10.44.2.10                 ethernet1/3
admin@PA-200>
```

## Clear Leases Before They Expire Automatically

The following example shows how to release expired DHCP Leases of an interface (server) before the hold timer releases them automatically. Those addresses will be available in the IP pool again.

```
admin@PA-200> clear dhcp lease interface ethernet1/2 expired-only
```

The following example shows how to release the lease of a particular IP address:

```
admin@PA-200> clear dhcp lease interface ethernet1/2 ip 192.168.3.1
```

The following example shows how to release the lease of a particular MAC address:

```
admin@PA-200> clear dhcp lease interface ethernet1/2 mac f0:2c:ae:29:71:34
```

© Palo Alto Networks, Inc.

## View DHCP Client Information

To view the status of IP address leases sent to the firewall when it is acting as a DHCP client, use the `show dhcp client state` *`interface_name`* command or the following command:

```
admin@PA-200> show dhcp client state all
Interface       State        IP              Gateway         Leased-until
--------------------------------------------------------------------------
ethernet1/1     Bound        10.43.14.80     10.43.14.1      70315
admin@PA-200>
```

## Gather Debug Output about DHCP

To gather debug output about DHCP, use one of the following commands:

```
admin@PA-200> debug dhcpd
admin@PA-200> debug management-server dhcpd
```

 © Palo Alto Networks, Inc.

# NAT

This section describes Network Address Translation (NAT) and how to configure NAT rules and features. NAT was introduced to solve the problem of an organization not having enough public, globally-routable IPv4 addresses assigned to it by the Internet Assigned Numbers Authority (IANA) or a Regional Internet Registry. NAT translates private, non-routable IPv4 addresses to one or more globally-routable IPv4 addresses, thereby conserving an organization's routable IP addresses.

Since its origination, the use of NAT has expanded. For example, NAT is used as a way to not disclose the real IP addresses of hosts that need access to public addresses. It is also used to manage traffic by performing port forwarding. NAT can be used to solve network design challenges, enabling networks with identical IP subnets to communicate with each other.

The NAT64 option translates between IPv6 and IPv4 addresses, providing connectivity between networks using disparate IP addressing schemes, and therefore a migration path to IPv6 addressing. IPv6-to-IPv6 Network Prefix Translation (NPTv6) translates one IPv6 prefix to another IPv6 prefix. PAN-OS supports all of these functions.

If you use private IP addresses within your internal networks, you must use NAT to translate the private addresses to public addresses that can be routed on external networks. In PAN-OS, you create NAT policy rules that instruct the firewall which packet addresses need translation and what the translated addresses are.


- NAT Policy Rules
- Source NAT and Destination NAT
- NAT Rule Capacities
- Dynamic IP and Port NAT Oversubscription
- Dataplane NAT Memory Statistics
- Configure NAT
- NAT Configuration Examples

# NAT Policy Rules


- NAT Policy Rule Functionality
- NAT Address Pools Identified as Address Objects
- Proxy ARP for NAT Address Pools


## NAT Policy Rule Functionality

You configure a NAT rule to match a packet's source zone and destination zone, at a minimum. In addition to zones, you can configure matching criteria based on the packet's destination interface, source and destination address, and service. You can configure multiple NAT rules. The firewall evaluates the rules in order from the top down. Once a packet matches the criteria of a single NAT rule, the packet is not subjected to additional NAT rules. Therefore, your list of NAT rules should be in order from most specific to least specific so that packets are subjected to the most specific rule you created for them.

Static NAT rules do not have precedence over other forms of NAT. Therefore, for static NAT to work, the static NAT rules must be above all other NAT rules in the list on the firewall.

NAT rules provide address translation, and are different from security policy rules, which allow or deny packets. It is important to understand the firewall's flow logic when it applies NAT rules and security policy rules so that you can determine what rules you need, based on the zones you have defined.

Upon ingress, the firewall inspects the packet and does a route lookup to determine the egress interface and zone. Then the firewall determines if the packet matches one of the NAT rules that have been defined, based on source and/or destination zone. It then evaluates and applies any security policies that match the packet based on the original (pre-NAT) source and destination addresses, but the post-NAT zones. Finally, upon egress, for a matching NAT rule, the firewall translates the source and/or destination address and port numbers.

Keep in mind that the translation of the IP address and port do not occur until the packet leaves the firewall. The NAT rules and security policies apply to the original IP address (the pre-NAT address). A NAT rule is configured based on the zone associated with a pre-NAT IP address.

Security policies differ from NAT rules because security policies examine post-NAT zones to determine whether the packet is allowed or not. Because the very nature of NAT is to modify source or destination IP addresses, which can result in modifying the packet's outgoing interface and zone, security policies are enforced on the post-NAT zone.

You can verify the NAT rules processed by using the CLI `test nat-policy-match` command in operational mode. For example:

```
user@device1> test nat-policy-match ?
+ destination       Destination IP address
+ destination-port  Destination port
+ from              From zone
+ ha-device-id      HA Active/Active device ID
+ protocol          IP protocol value
+ source            Source IP address
+ source-port       Source port
+ to                To Zone
+ to-interface      Egress interface to use
  |                 Pipe through a command
```

 © Palo Alto Networks, Inc.

```
<Enter>              Finish input

user@device1> test nat-policy-match from l3-untrust source 10.1.1.1 destination
66.151.149.20 destination-port 443 protocol 6

Destination-NAT: Rule matched: CA2-DEMO
66.151.149.20:443 => 192.168.100.15:443
```

## NAT Address Pools Identified as Address Objects

When configuring a **Dynamic IP** or **Dynamic IP and Port** NAT address pool in a NAT policy rule, it is typical to configure the pool of translated addresses with address objects. Each address object can be a host IP address, IP address range, or IP subnet.

Because both NAT rules and security policy rules use address objects, it is a best practice to distinguish between them by naming an address object used for NAT with a prefix, such as "NAT-*name*."

## Proxy ARP for NAT Address Pools

NAT address pools are not bound to any interfaces. The following figure illustrates the behavior of the firewall when it is performing proxy ARP for an address in a NAT address pool.

The firewall performs source NAT for a client, translating the source address 1.1.1.1 to the address in the NAT pool, 2.2.2.2. The translated packet is sent on to a router.

For the return traffic, the router does not know how to reach 2.2.2.2 (because the IP address 2.2.2.2 is just an address in the NAT address pool), so it sends an ARP request packet to the firewall.

- If the address pool (2.2.2.2) is in the same subnet as the egress/ingress interface IP address (2.2.2.3/24), the firewall can send a proxy ARP reply to the router, indicating the Layer 2 MAC address of the IP address, as shown in the figure above.
- If the address pool (2.2.2.2) is not a subnet of an interface on the firewall, the firewall will not send a proxy ARP reply to the router. This means that the router must be configured with the necessary route to know where to send packets destined for 2.2.2.2, in order to ensure the return traffic is routed back to the firewall, as shown in the figure below.

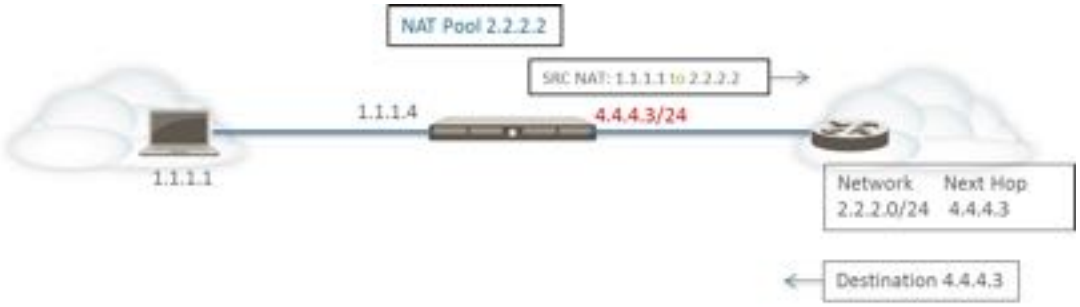


 © Palo Alto Networks, Inc.

# Source NAT and Destination NAT

The firewall supports both source address and/or port translation and destination address and/or port translation.

## Source NAT

Source NAT is typically used by internal users to access the Internet; the source address is translated and thereby kept private. There are three types of source NAT:

- **Dynamic IP and Port (DIPP)**—Allows multiple hosts to have their source IP addresses translated to the same public IP address with different port numbers. The dynamic translation is to the next available address in the NAT address pool, which you configure as a **Translated Address** pool be to an IP address, range of addresses, a subnet, or a combination of these.

  As an alternative to using the next address in the NAT address pool, DIPP allows you to specify the address of the **Interface** itself. The advantage of specifying the interface in the NAT rule is that the NAT rule will be automatically updated to use any address subsequently acquired by the interface. DIPP is sometimes referred to as interface-based NAT or network address port translation (NAPT).

  DIPP has a default NAT oversubscription rate, which is the number of times that the same translated IP address and port pair can be used concurrently. For more information, see Dynamic IP and Port NAT Oversubscription and Modify the Oversubscription Rate for DIPP NAT.

- **Dynamic IP**—Allows the one-to-one, dynamic translation of a source IP address only (no port number) to the next available address in the NAT address pool. The size of the NAT pool should be equal to the number of internal hosts that require address translations. By default, if the source address pool is larger than the NAT address pool and eventually all of the NAT addresses are allocated, new connections that need address translation are dropped. To override this default behavior, use **Advanced (Dynamic IP/Port Fallback)** to enable use of DIPP addresses when necessary. In either event, as sessions terminate and the addresses in the pool become available, they can be allocated to translate new connections.

  Dynamic IP NAT supports the option for you to Reserve Dynamic IP NAT Addresses.

- **Static IP**—Allows the 1-to-1, static translation of a source IP address, but leaves the source port unchanged. A common scenario for a static IP translation is an internal server that must be available to the Internet.

## Destination NAT

Destination NAT is performed on incoming packets, when the firewall translates a public destination address to a private address. Destination NAT does not use address pools or ranges. It is a 1-to-1, static translation with the option to perform port forwarding or port translation.

- **Static IP**—Allows the 1-to-1, static translation of a destination IP address and optionally the port number.

One common use of destination NAT is to configure several NAT rules that map a single public destination address to several private destination host addresses assigned to servers or services. In this case, the destination port numbers are used to identify the destination hosts. For example:

© Palo Alto Networks, Inc.

- **Port Forwarding**—Can translate a public destination address and port number to a private destination address, but keeps the same port number.
- **Port Translation**—Can translate a public destination address and port number to a private destination address and a different port number, thus keeping the real port number private. It is configured by entering a **Translated Port** on the **Translated Packet** tab in the NAT policy rule. See the Destination NAT with Port Translation Example.

 © Palo Alto Networks, Inc.

## NAT Rule Capacities

The number of NAT rules allowed is based on the firewall platform. Individual rule limits are set for static, Dynamic IP (DIP), and Dynamic IP and Port (DIPP) NAT. The sum of the number of rules used for these NAT types cannot exceed the total NAT rule capacity. For DIPP, the rule limit is based on the device's oversubscription setting (8, 4, 2, or 1) and the assumption of one translated IP address per rule. To see platform-specific NAT rule limits and translated IP address limits, use the Compare Firewalls tool.

Consider the following when working with NAT rules:

- If you run out of pool resources, you cannot create more NAT rules, even if the platform's maximum rule count has not been reached.
- If you consolidate NAT rules, the logging and reporting will also be consolidated. The statistics are provided per the rule, not per all of the addresses within the rule. If you need granular logging and reporting, do not combine the rules.

# Dynamic IP and Port NAT Oversubscription

Dynamic IP and Port (DIPP) NAT allows you to use each translated IP address and port pair multiple times (8, 4, or 2 times) in concurrent sessions. This reusability of an IP address and port (known as oversubscription) provides scalability for customers who have too few public IP addresses. The design is based on the assumption that hosts are connecting to different destinations, therefore sessions can be uniquely identified and collisions are unlikely. The oversubscription rate in effect multiplies the original size of the address/port pool to 8, 4, or 2 times the size. For example, the default limit of 64K concurrent sessions allowed, when multiplied by an oversubscription rate of 8, results in 512K concurrent sessions allowed.

The oversubscription rates that are allowed vary based on the platform. The oversubscription rate is global; it applies to the device. This oversubscription rate is set by default and consumes memory, even if you have enough public IP addresses available to make oversubscription unnecessary. You can reduce the rate from the default setting to a lower setting or even 1 (which means no oversubscription). By configuring a reduced rate, you decrease the number of source device translations possible, but increase the DIP and DIPP NAT rule capacities. To change the default rate, see Modify the Oversubscription Rate for DIPP NAT.

If you select **Platform Default**, your explicit configuration of oversubscription is turned off and the default oversubscription rate for the platform applies, as shown in the table below. The **Platform Default** setting allows for an upgrade or downgrade of a software release.

The following table lists the default (highest) oversubscription rate for each platform.

| Platform | Default Oversubscription Rate |
|---|---|
| PA-200 | 2 |
| PA-500 | 2 |
| PA-2020 | 2 |
| PA-2050 | 2 |
| PA-3020 | 2 |
| PA-3050 | 2 |
| PA-3060 | 2 |
| PA-4020 | 4 |
| PA-4050 | 8 |
| PA-4060 | 8 |
| PA-5020 | 4 |
| PA-5050 | 8 |
| PA-5060 | 8 |
| PA-7050 | 8 |
| PA-7080 | 8 |
| VM-100 | 1 |

 © Palo Alto Networks, Inc.

| Platform | Default Oversubscription Rate |
|---|---|
| VM-200 | 1 |
| VM-300 | 2 |
| VM-1000-HV | 2 |

The firewall supports a maximum of 256 translated IP addresses per NAT rule, and each platform supports a maximum number of translated IP addresses (for all NAT rules combined). If oversubscription causes the maximum translated addresses per rule (256) to be exceeded, the firewall will automatically reduce the oversubscription ratio in an effort to have the commit succeed. However, if your NAT rules result in translations that exceed the maximum translated addresses for the platform, the commit will fail.

© Palo Alto Networks, Inc.

## Dataplane NAT Memory Statistics

The `show running global-ippool` command displays statistics related to NAT memory consumption for a pool. The Size column displays the number of bytes of memory that the resource pool is using. The Ratio column displays the oversubscription ratio (for DIPP pools only). The lines of pool and memory statistics are explained in the following sample output:

For NAT pool statistics for a virtual system, the `show running ippool` command has columns indicating the memory size used per NAT rule and the oversubscription ratio used (for DIPP rules). The following is sample output for the command.

A field in the output of the `show running nat-rule-ippool rule` command shows the memory (bytes) used per NAT rule. The following is sample output for the command, with the memory usage for the rule encircled.

 © Palo Alto Networks, Inc.

## Configure NAT

Perform the following tasks to configure various aspects of NAT. In addition to the examples below, there are examples in the section NAT Configuration Examples.

- Translate Internal Client IP Addresses to Your Public IP Address (Source DIPP NAT)
- Enable Clients on the Internal Network to Access your Public Servers (Destination U-Turn NAT)
- Enable Bi-Directional Address Translation for Your Public-Facing Servers (Static Source NAT)
- Modify the Oversubscription Rate for DIPP NAT
- Disable NAT for a Specific Host or Interface
- Reserve Dynamic IP NAT Addresses

The NAT example in this section is based on the following topology, which was also used in Getting Started for setting up interfaces and zones:

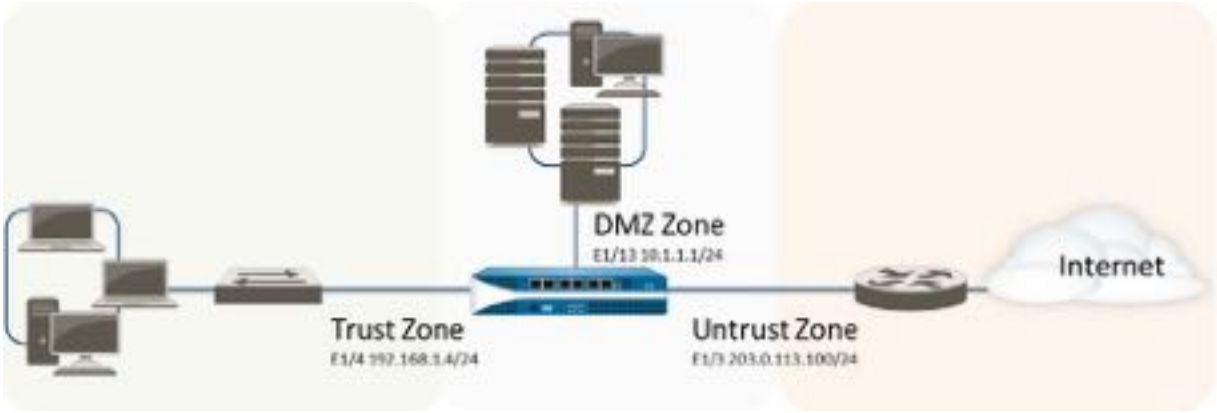


Based on the topology initially used in Getting Started to create the interfaces and zones, there are three NAT policies we need to create as follows:

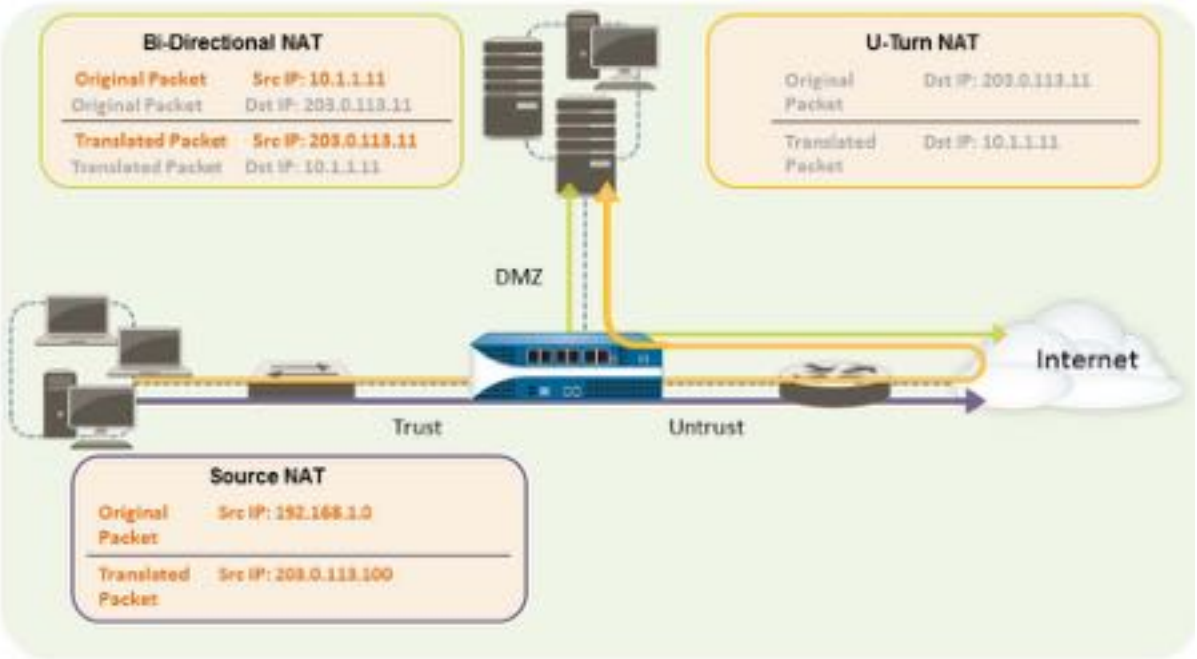


- To enable the clients on the internal network to access resources on the Internet, the internal 192.168.1.0 addresses will need to be translated to publicly routable addresses. In this case, we will configure source NAT (the purple enclosure and arrow above), using the egress interface address, 203.0.113.100, as the source address in all packets that leave the firewall from the internal zone. See Translate Internal Client IP Addresses to Your Public IP Address (Source DIPP NAT) for instructions.
- To enable clients on the internal network to access the public web server in the DMZ zone, we must configure a NAT rule that redirects the packet from the external network, where the original routing table lookup will determine it should go based on the destination address of 203.0.113.11 within the packet, to the actual address of the web server on the DMZ network of 10.1.1.11. To do this you must create a NAT rule from the trust zone (where the source address in the packet is) to the untrust zone (where the original destination address is) to translate the destination address to an address in the DMZ zone. This type of destination NAT is called *U-Turn NAT* (the yellow enclosure and arrow above). See Enable Clients on the Internal Network to Access your Public Servers (Destination U-Turn NAT) for instructions.
- To enable the web server—which has both a private IP address on the DMZ network and a public-facing address for access by external users—to both send and receive requests, the firewall must translate the incoming packets from the public IP address to the private IP address and the outgoing packets from the private IP address to the public IP address. On the firewall, you can accomplish this with a single bi-directional static source NAT policy (the green enclosure and arrow above). See Enable Bi-Directional Address Translation for Your Public-Facing Servers (Static Source NAT).

 © Palo Alto Networks, Inc.

## Translate Internal Client IP Addresses to Your Public IP Address (Source DIPP NAT)

When a client on your internal network sends a request, the source address in the packet contains the IP address for the client on your internal network. If you use private IP address ranges internally, the packets from the client will not be able to be routed on the Internet unless you translate the source IP address in the packets leaving the network into a publicly routable address.

On the firewall you can do this by configuring a source NAT policy that translates the source address (and optionally the port) into a public address. One way to do this is to translate the source address for all packets to the egress interface on your firewall, as shown in the following procedure.

| **Configure Source NAT** | |
|---|---|
| Step 1 Create an address object for the external IP address you plan to use. | 1. Select **Objects > Addresses** and then click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the object.<br>3. Select **IP Netmask** from the **Type** drop-down and then enter the IP address of the external interface on the firewall, 203.0.113.100 in this example.<br> <br>4. To save the address object, click **OK**.<br>Although you do not have to use address objects in your policies, it is a best practice because it simplifies administration by allowing you to make updates in one place rather than having to update every policy where the address is referenced. |

| Configure Source NAT (Continued) | |
|---|---|
| Step 2 Create the NAT policy. | 1. Select **Policies > NAT** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the policy.<br>3. (Optional) Enter a tag, which is a keyword or phrase that allows you to sort or filter policies.<br>4. For **NAT Type**, select **ipv4** (default).<br>5. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.<br> <br>6. On the **Translated Packet** tab, select **Dynamic IP And Port** from the **Translation Type** drop-down in the Source Address Translation section of the screen.<br>7. For **Address Type**, there are two choices. You could select **Translated Address** and then click **Add**. Select the address object you just created.<br> 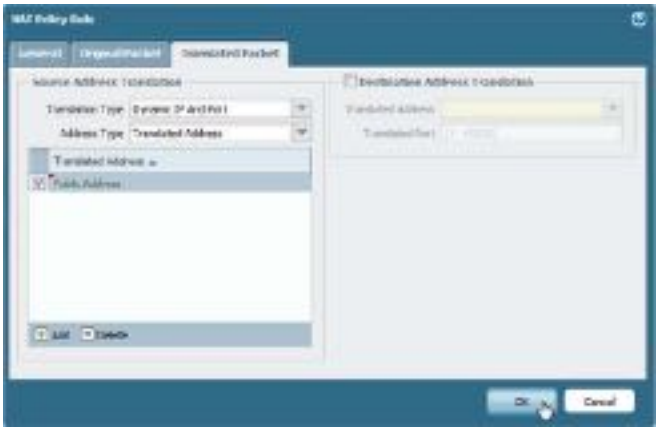 <br>An alternative **Address Type** is **Interface Address**, in which case the translated address will be the IP address of the interface. For this choice, you would select an **Interface** and optionally an **IP Address** if the interface has more than one IP address.<br>8. Click **OK** to save the NAT policy. |
| Step 3 Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

| Configure Source NAT (Continued) | |
|---|---|
| Step 4 (Optional) Access the CLI to verify the translation. | 1. Use the `show session all` command to view the session table, where you can verify the source IP address and port and the corresponding translated IP address and port.<br>2. Use the `show session id <id_number>` to view more details about a session.<br>3. If you configured Dynamic IP NAT, use the `show counter global filter aspect session severity drop \| match nat` command to see if any sessions failed due to NAT IP allocation. If all of the addresses in the Dynamic IP NAT pool are allocated when a new connection is supposed to be translated, the packet will be dropped. |

## Enable Clients on the Internal Network to Access your Public Servers (Destination U-Turn NAT)

When a user on the internal network sends a request for access to the corporate web server in the DMZ, the DNS server will resolve it to the public IP address. When processing the request, the firewall will use the original destination in the packet (the public IP address) and route the packet to the egress interface for the untrust zone. In order for the firewall to know that it must translate the public IP address of the web server to an address on the DMZ network when it receives requests from users on the trust zone, you must create a destination NAT rule that will enable the firewall to send the request to the egress interface for the DMZ zone as follows.

| Configure U-Turn NAT | |
|---|---|
| Step 1 Create an address object for the web server. | 1. Select **Objects > Addresses** and click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the object.<br>3. Select **IP Netmask** from the **Type** drop-down and enter the public IP address of the web server, 203.0.113.11 in this example.<br> <br>4. Click **OK**. |

| Configure U-Turn NAT (Continued) | |
|---|---|
| Step 1 Create the NAT policy. | 1. Select **Policies > NAT** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the NAT rule.<br>3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.<br>4. In the **Destination Address** section, click **Add** and select the address object you created for your public web server.<br> 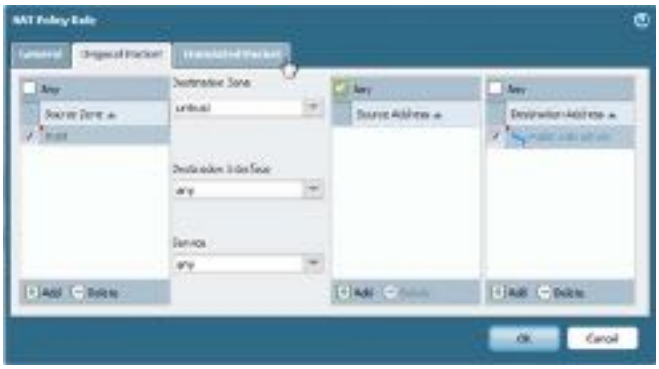 <br>5. On the **Translated Packet** tab, select the **Destination Address Translation** check box and then enter the IP address that is assigned to the web server interface on the DMZ network, 10.1.1.11 in this example.<br> <br>6. Click **OK** to save the NAT policy. |
| Step 2 Save the configuration. | Click **Commit**. |

## Enable Bi-Directional Address Translation for Your Public-Facing Servers (Static Source NAT)

When your public-facing servers have private IP addresses assigned on the network segment where they are physically located, you need a source NAT rule to translate the source address of the server to the external address upon egress. You create a static NAT rule to translate the internal source address, 10.1.1.11, to the external web server address, 203.0.113.11 in our example.

However, a public-facing server must be able to both send and receive packets. You need a reciprocal policy that translates the public address (the destination IP address in incoming packets from Internet users) into the private address so that the firewall can route the packet to your DMZ network. You create a bi-directional static NAT rule, as described in the following procedure. Bi-directional translation is an option for static NAT only.

 © Palo Alto Networks, Inc.

| Configure Bi-Directional NAT | |
|---|---|
| Step 1 Create an address object for the web server's internal IP address. | 1. Select **Objects > Addresses** and click **Add**.<br>2. Enter a **Name** and optionally a **Description** for the object.<br>3. Select **IP Netmask** from the **Type** drop-down and enter the IP address of the web server on the DMZ network, 10.1.1.11 in this example.<br> <br>4. Click **OK**.<br>If you did not already create an address object for the *public* address of your web server, you should create that object now. |

| Configure Bi-Directional NAT (Continued) | |
|---|---|
| Step 2 Create the NAT policy. | 1. Select **Policies > NAT** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the NAT rule.<br>3. On the **Original Packet** tab, select the zone you created for your DMZ in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.<br>4. In the **Source Address** section, click **Add** and select the address object you created for your internal web server address.<br> 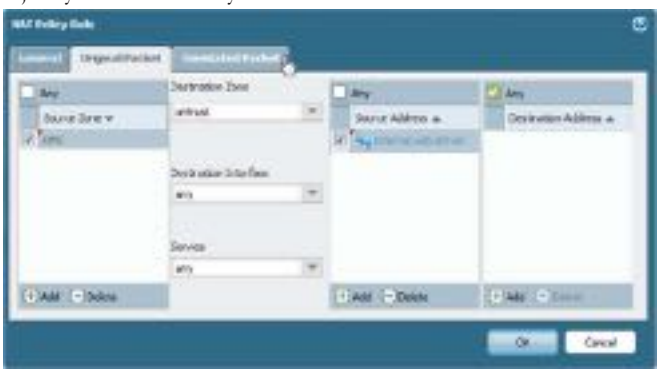 <br>5. On the **Translated Packet** tab, select **Static IP** from the **Translation Type** drop-down in the **Source Address Translation** section and then select the address object you created for your external web server address from the **Translated Address** drop-down.<br>6. In the **Bi-directional** field, select **Yes**.<br> <br>7. Click **OK** to save the NAT policy. |
| Step 3 Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

## Modify the Oversubscription Rate for DIPP NAT

If you have enough public IP addresses that you do not need to use DIPP NAT oversubscription, you can reduce the oversubscription rate and thereby gain more DIP and DIPP NAT rules allowed.

| **Set NAT Oversubscription** | | |
|---|---|---|
| Step 1 | View the DIPP NAT oversubscription rate. | 1. Select **Device > Setup > Session > Session Settings.** View the **NAT Oversubscription Rate** setting. |
| Step 2 | Set the DIPP NAT oversubscription rate. | 1. Click the Edit icon in the Session Settings section.<br>2. In the **NAT Oversubscription Rate** drop-down, select **1x**, **2x**, **4x**, or **8x**, depending on which ratio you want.<br>The **Platform Default** setting applies the default oversubscription setting for the platform. If you want no oversubscription, select **1x**.<br>3. Click **OK** and **Commit** the change. |

## Disable NAT for a Specific Host or Interface

Both source NAT and destination NAT rules can be configured to disable address translation. You may have exceptions where you do not want NAT to occur for a certain host in a subnet or for traffic exiting a specific interface. The following procedure shows how to disable source NAT for a host.

| **Create a Source NAT Exemption** | | |
|---|---|---|
| Step 1 | Create the NAT policy. | 1. Select **Policies > NAT** and click **Add**.<br>2. Enter a descriptive **Name** for the policy.<br>3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.<br>4. For **Source Address**, click **Add** and enter the host address. Click **OK**.<br>5. On the **Translated Packet** tab, select **None** from the **Translation Type** drop-down in the Source Address Translation section of the screen.<br>6. Click **OK** to save the NAT policy. |
| Step 2 | Save the configuration. | Click **Commit**. |

NAT rules are processed in order from the top to the bottom, so place the NAT exemption policy before other NAT policies to ensure it is processed before an address translation occurs for the sources you want to exempt.

## Reserve Dynamic IP NAT Addresses

You can reserve Dynamic IP NAT addresses (for a configurable period of time) to prevent them from being allocated as translated addresses to a different source IP address that needs translation. When configured, the reservation applies to all of the translated Dynamic IP addresses in progress and any new translations.

For both translations in progress and new translations, when a source IP address is translated to an available translated IP address, that pairing is retained even after all sessions related to that specific source IP are expired. The reservation timer for each source IP address begins after all sessions that use that source IP address translation expire. Dynamic IP NAT is a one-to-one translation; one source IP address translates to one translated IP address that is chosen dynamically from those addresses available in the configured pool. Therefore, a translated IP address that is reserved is not available for any other source IP address until the reservation expires because a new session has not started. The timer is reset each time a new session for a source IP/translated IP mapping begins, after a period when no sessions were active.

By default, no addresses are reserved. You can reserve Dynamic IP NAT addresses for the firewall or for a virtual system.

| **Reserve Dynamic IP NAT Addresses for a Firewall** | |
|---|---|
| Step 1 | `user@device1# set setting nat reserve-ip yes` |
| Step 2 | `user@device1# set setting nat reserve-time <1-604800 secs>` |

| **Reserve Dynamic IP NAT Addresses for a Virtual System** | |
|---|---|
| Step 1 | `user@device1# set vsys <vsysid> setting nat reserve-ip yes` |
| Step 2 | `user@device1# set vsys <vsysid> setting nat reserve-time <1-604800 secs>` |

For example, suppose there is a Dynamic IP NAT pool of 30 addresses and there are 20 translations in progress when the `nat reserve-time` is set to 28800 seconds (8 hours). Those 20 translations are now reserved, so that when the last session (of any application) that uses each source IP/translated IP mapping expires, the translated IP address is reserved for only that source IP address for 8 hours, in case that source IP address needs translation again. Additionally, as the 10 remaining translated addresses are allocated, they each are reserved for their source IP address, each with a timer that begins when the last session for that source IP address expires.

In this manner, each source IP address can be repeatedly translated to its same NAT address from the pool; another host will not be assigned a reserved translated IP address from the pool, even if there are no active sessions for that translated address.

Suppose a source IP/translated IP mapping has all of its sessions expire, and the reservation timer of 8 hours begins. After a new session for that translation begins, the timer stops, and the sessions continue until they all end, at which point the reservation timer starts again, reserving the translated address.

The reservation timer remain in effect on the Dynamic IP NAT pool until you disable it by entering the `set setting nat reserve-ip no` command or you change the `nat reserve-time` to a different value.

The CLI commands for reservations do not affect Dynamic IP and Port (DIPP) or Static IP NAT pools.

 © Palo Alto Networks, Inc.

# NAT Configuration Examples


- Destination NAT Example—One-to-One Mapping
- Destination NAT with Port Translation Example
- Destination NAT Example—One-to-Many Mapping
- Source and Destination NAT Example
- Virtual Wire Source NAT Example
- Virtual Wire Static NAT Example
- Virtual Wire Destination NAT Example


## Destination NAT Example—One-to-One Mapping

The most common mistakes when configuring NAT and security rules are the references to the zones and address objects. The addresses used in destination NAT rules always refer to the original IP address in the packet (that is, the pre-translated address). The destination zone in the NAT rule is determined after the route lookup of the destination IP address in the original packet (that is, the pre-NAT destination IP address).

The addresses in the security policy also refer to the IP address in the original packet (that is, the pre-NAT address). However, the destination zone is the zone where the end host is physically connected. In other words, the destination zone in the security rule is determined after the route lookup of the post-NAT destination IP address.

In the following example of a one-to-one destination NAT mapping, users from the zone named untrust-l3 access the server 10.1.1.100 in the zone named dmz-l3 using the IP address 1.1.1.100.

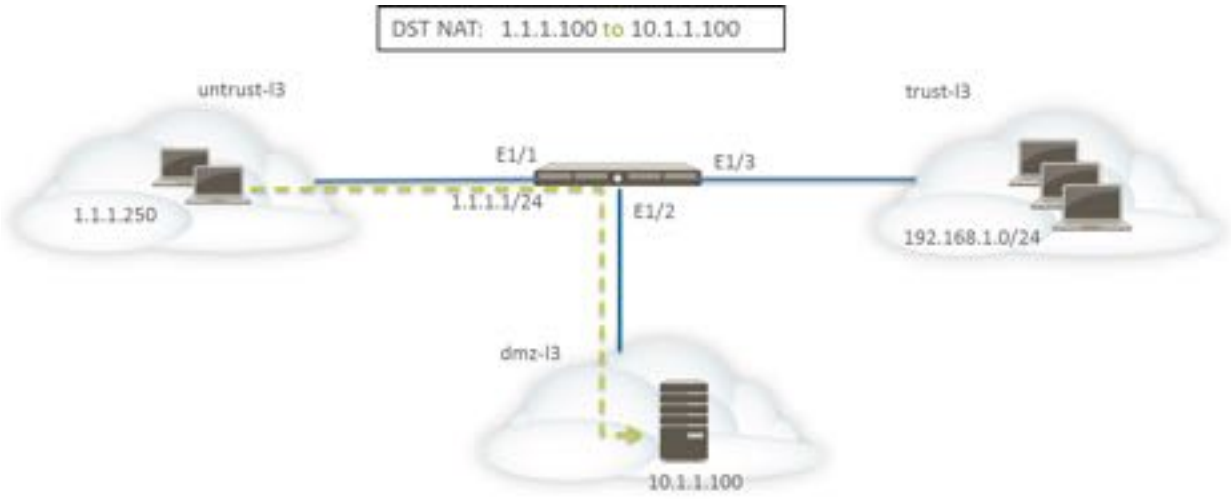


Before configuring the NAT rules, consider the sequence of events for this scenario.

- Host 1.1.1.250 sends an ARP request for the address 1.1.1.100 (the public address of the destination server).

- The firewall receives the ARP request packet for destination 1.1.1.100 on the Ethernet1/1 interface and processes the request. The firewall responds to the ARP request with its own MAC address because of the destination NAT rule configured.
- The NAT rules are evaluated for a match. For the destination IP address to be translated, a destination NAT rule from zone untrust-l3 to zone untrust-l3 must be created to translate the destination IP of 1.1.1.100 to 10.1.1.100.
- After determining the translated address, the firewall performs a route lookup for destination 10.1.1.100 to determine the egress interface. In this example, the egress interface is Ethernet1/2 in zone dmz-l3.
- The firewall performs a security policy lookup to see if the traffic is permitted from zone untrust-l3 to dmz-l3.

The direction of the policy matches the ingress zone and the zone where the server is physically located.

The security policy refers to the IP address in the original packet, which has a destination address of 1.1.1.100.

- The firewall forwards the packet to the server out egress interface Ethernet1/2. The destination address is changed to 10.1.1.100 as the packet leaves the firewall.

For this example, the following address objects are configured:

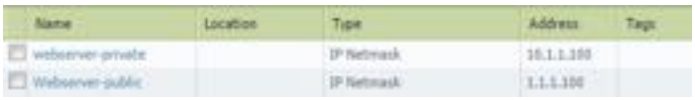


The configured NAT rule would look like this:

The direction of the NAT rules is based on the result of route lookup.

The configured security policy to provide access to the server from the untrust-l3 zone would look like this:

## Destination NAT with Port Translation Example

In this example, the web server is configured to listen for HTTP traffic on port 8080. The clients access the web server using the IP address 1.1.1.100 and TCP Port 80. The destination NAT rule is configured to translate both IP address and port to 10.1.1.100 and TCP port 8080.

 © Palo Alto Networks, Inc.

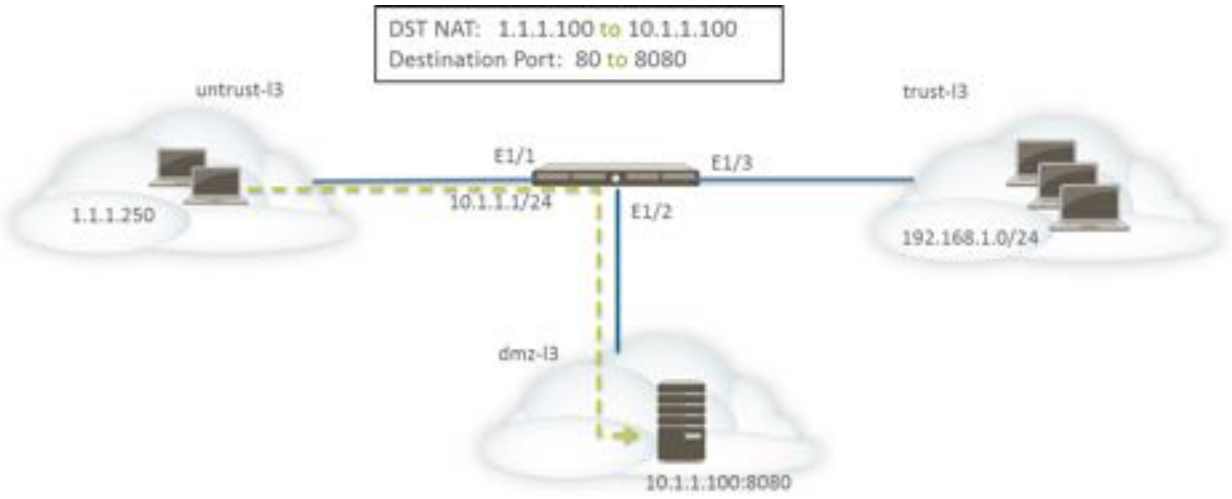


The following NAT and security rules must be configured on the firewall:

| | Original Packet | | | | | | Translated Packet | |
|---|---|---|---|---|---|---|---|---|
| Name | Source Zone | Destination Zone | Destination Interface | Source Address | Destination Address | Service | Source Translation | Destination Translation |
| Dst NAT-webserv | untrust-l3 | untrust-l3 | none | any | Servers-public | service-http | none | address: webserver-private<br>port: 8080 |

| | Source | | Destination | | | |
|---|---|---|---|---|---|---|
| Name | Zone | Address | Zone | Address | Application | Service |
| Webserver access | untrust-l3 | any | dmz-l3 | Servers-public | web-browsing | any |

Use the `show session all` CLI command to verify the translation.

## Destination NAT Example—One-to-Many Mapping

In this example, one IP address maps to two different internal hosts. The firewall uses the application to identify the internal host to which the firewall forwards the traffic.

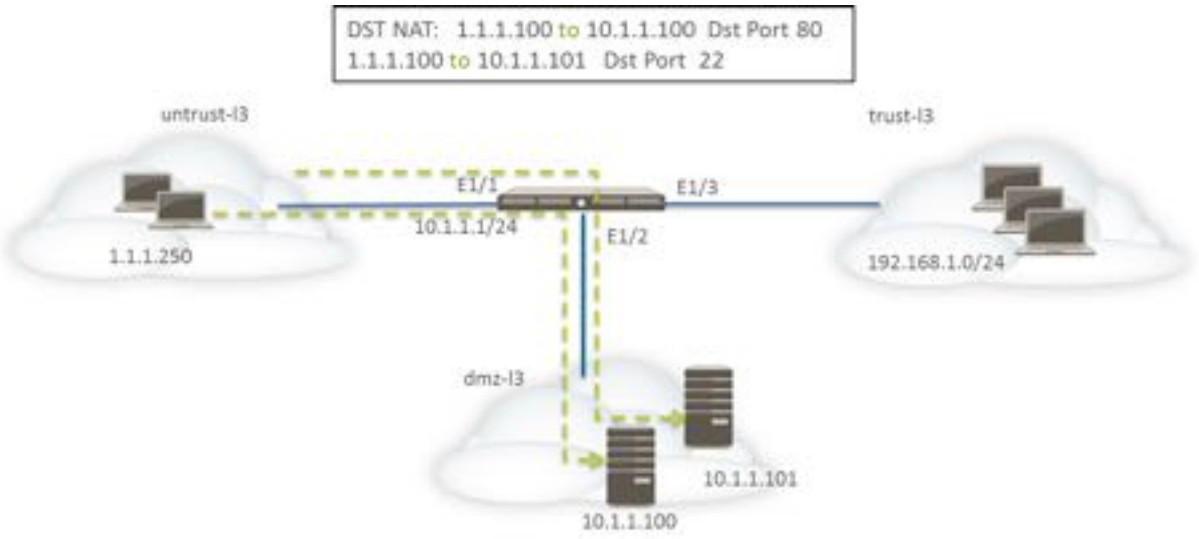


All HTTP traffic is sent to host 10.1.1.100 and SSH traffic is sent to server 10.1.1.101. The following address objects are required:

- Address object for the one pre-translated IP address of the server
- Address object for the real IP address of the SSH server
- Address object for the real IP address of the web server

The configured address objects would look like this:

| Name | Type | Address |
|---|---|---|
| Servers-public | IP Netmask | 1.1.1.100 |
| SSH-server | IP Netmask | 10.1.1.101 |
| webserver-private | IP Netmask | 10.1.1.100 |

The NAT rules would look like this:

| | Original Packet | | | | | | Translated Packet | |
|---|---|---|---|---|---|---|---|---|
| Name | Source Zone | Destination Zone | Destination Interface | Source Address | Destination Address | Service | Source Translation | Destination Translation |
| Dst NAT-webserv | untrust-l3 | untrust-l3 | none | any | Servers-public | service-http | none | address: webserver-private |
| Dst NAT-SSH | untrust-l3 | untrust-l3 | none | any | Servers-public | custom-ssh | none | address: SSH-server |

The security rules would look like this:

| | Source | | Destination | | | | |
|---|---|---|---|---|---|---|---|
| Name | Zone | Address | Zone | Address | Application | Service | Action |
| Webserver access | untrust-l3 | any | dmz-l3 | Servers-public | web-browsing | application-default | ✓ |
| SSH access | untrust-l3 | any | dmz-l3 | Servers-public | ssh | application-default | ✓ |

 © Palo Alto Networks, Inc.

## Source and Destination NAT Example

In this example, NAT rules translate both the source and destination IP address of packets between the clients and the server.

- Source NAT—The source addresses in the packets from the clients in the l3-trust zone to the server in the l3-untrust zone are translated from the private addresses in the network 192.168.1.0/24 to the IP address of the egress interface on the firewall (10.16.1.103). Dynamic IP and Port translation causes the port numbers to be translated also.
- Destination NAT—The destination addresses in the packets from the clients to the server are translated from the server's public address (80.80.80.80) to the server's private address (10.2.133.15).

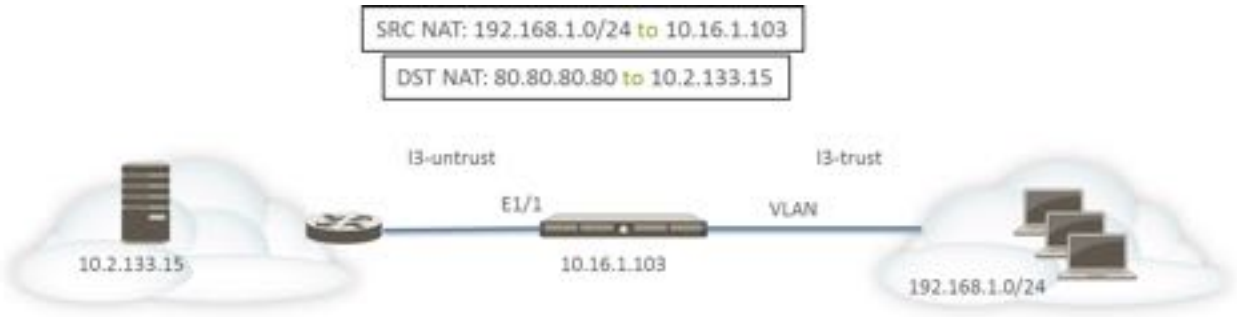


The following address objects are created for destination NAT.

- Server-Pre-NAT: 80.80.80.80
- Server-post-NAT: 10.2.133.15

The following screen shots illustrate how to configure the source and destination NAT policies for the example.

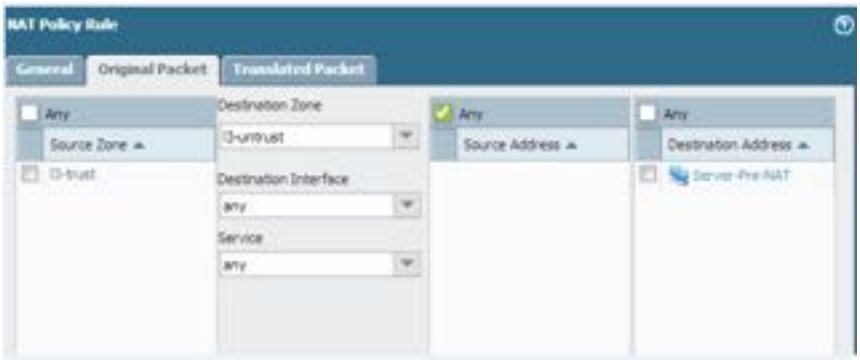

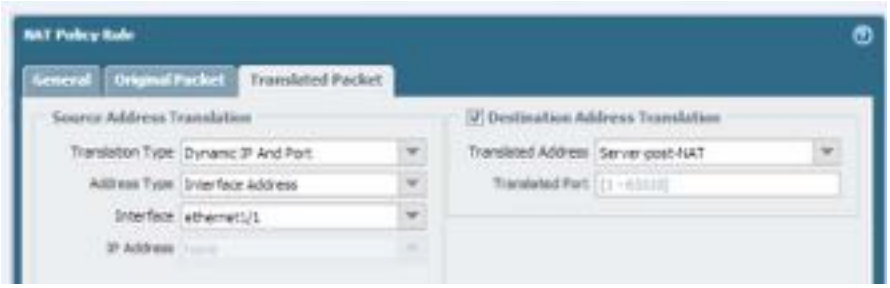


To verify the translations, use the CLI command `show session all filter destination 80.80.80.80`. Note that a client address 192.168.1.11 and its port number are translated to 10.16.1.103 and a port number. The destination address 80.80.80.80 is translated to 10.2.133.15.

## Virtual Wire Source NAT Example

Virtual wire deployment of a Palo Alto Networks firewall includes the benefit of providing security transparently to the end devices. It is possible to configure NAT for interfaces configured in a virtual wire. All of the NAT types are allowed: source NAT (Dynamic IP, Dynamic IP and Port, static) and destination NAT.

Because interfaces in a virtual wire do not have an IP address assigned, it is not possible to translate an IP address to an interface IP address. You must configure an IP address pool.

The firewall will not proxy ARP for NAT addresses. Ensure that routes are configured on the upstream and downstream devices. See Proxy ARP for NAT Address Pools for more explanation about proxy ARP.

Proper routing must be configured on the upstream and downstream routers in order for the packets to be translated in virtual wire mode.

In the source NAT and static NAT examples below, security policies (not shown) are configured from the virtual wire zone named vw-trust to the zone named vw-untrust.

In the following topology, two routers are configured to provide connectivity between subnets 1.1.1.0/24 and 3.1.1.0/24. The link between the routers is configured in subnet 2.1.1.0/30. Static routing is configured on both routers to establish connectivity between the networks. Before the firewall is deployed in the environment, the topology and the routing table for each router look like this:

Route on R1:

| Destination | Next Hop |
|---|---|
| 3.1.1.0/24 | 2.1.1.2 |

 © Palo Alto Networks, Inc.

Route on R2:

| Destination | Next Hop |
|---|---|
| 1.1.1.0/24 | 2.1.1.1 |

Now the firewall is deployed in virtual wire mode between the two Layer 3 devices. All communications from clients in network 1.1.1.0/24 accessing servers in network 3.1.1.0/24 are translated to an IP address in the range 2.1.1.9-2.1.1.14. A NAT IP address pool with range 2.1.1.9-2.1.1.14 is configured on the firewall.

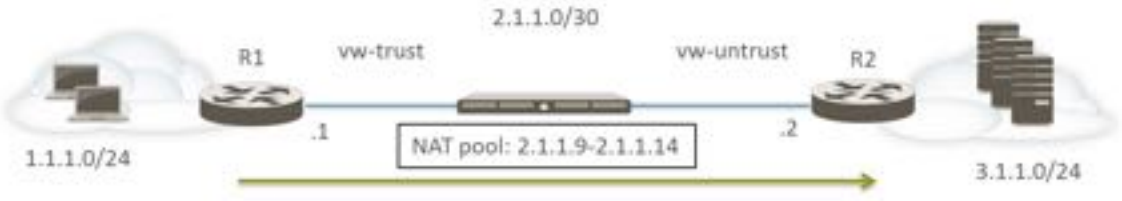


All connections from the clients in subnet 1.1.1.0/24 will arrive at router R2 with a translated source address in the range 2.1.1.9-2.1.1.14. The response from servers will be directed to these addresses. In order for source NAT to work, you must configure proper routing on router R2, so that packets destined for other addresses are not dropped. The routing table below shows the modified routing table on router R2. The route ensures the traffic to the destinations 2.1.1.9-2.1.1.14 (that is, hosts on subnet 2.1.1.8/29) will be sent back through the firewall to router R1.

Route on R2:

| Destination | Next Hop |
|---|---|
| 2.1.1.8/29 | 2.1.1.1 |

## Virtual Wire Static NAT Example

In this example, security policies are configured from the virtual wire zone named vw-trust to vw-untrust. Host 1.1.1.100 is statically translated to address 2.1.1.100. With the **Bi-directional** option enabled, the firewall generates a NAT policy from the vw-untrust zone to the vw-trust zone. Clients on the vw-untrust zone access the server using the IP address 2.1.1.100, which the firewall translates to 1.1.1.100. Any connections initiated by the server at 1.1.1.100 are translated to source IP address 2.1.1.100.

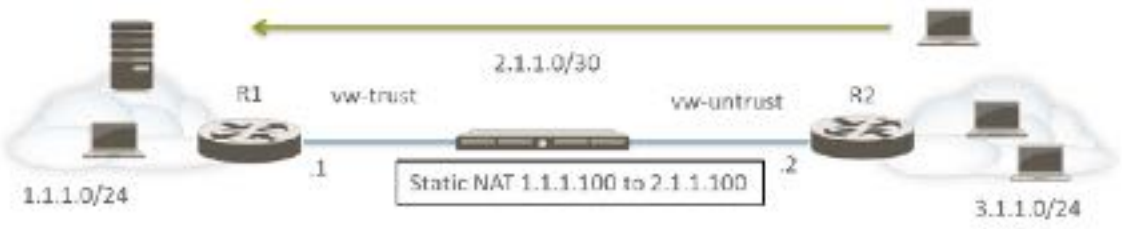

Route on R2:

| Destination | Next Hop |
|---|---|
| 2.1.1.100/32 | 2.1.1.1 |

## Virtual Wire Destination NAT Example

Clients in the vw-untrust zone access the server using the IP address 2.1.1.100, which the firewall translates to 1.1.1.100. Both the NAT and security policies must be configured from the vw-untrust zone to the vw-trust zone.

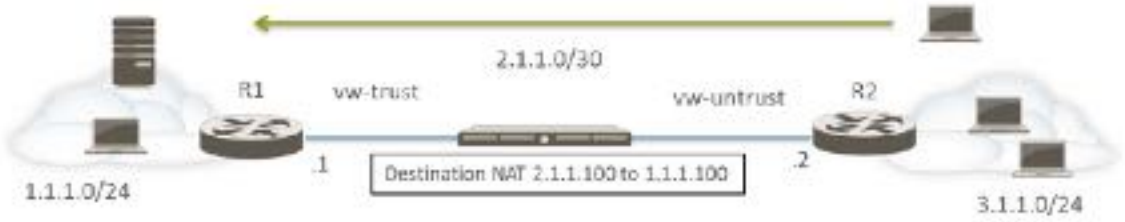


Route on R2:

| Destination | Next Hop |
|---|---|
| 2.1.1.100/32 | 2.1.1.1 |

 © Palo Alto Networks, Inc.

# NPTv6

IPv6-to-IPv6 Network Prefix Translation (NPTv6) performs a stateless, static translation of one IPv6 prefix to another IPv6 prefix (port numbers are not changed). There are four primary benefits of NPTv6:

- You can prevent the asymmetrical routing problems that result from Provider Independent addresses being advertised from multiple datacenters.
- NPTv6 allows more specific routes to be advertised so that return traffic arrives at the same firewall that transmitted the traffic.
- Private and public addresses are independent; you can change one without affecting the other.
- You have the ability to translate Unique Local Addresses to globally routable addresses.

This topic builds on a basic understanding of NAT. You should be sure you are familiar with NAT concepts before configuring NPTv6.

▲ NPTv6 Overview
▲ How NPTv6 Works
▲ NDP Proxy
▲ NPTv6 and NDP Proxy Example
▲ Create an NPTv6 Policy

# NPTv6 Overview

This section describes IPv6-to-IPv6 Network Prefix Translation (NPTv6) and how to configure it. NPTv6 is defined in RFC 6296. Palo Alto Networks does not implement all functionality defined in the RFC, but is compliant with the RFC in the functionality it has implemented.

NPTv6 performs stateless translation of one IPv6 prefix to another IPv6 prefix. It is stateless, meaning that it does not keep track of ports or sessions on the addresses translated. NPTv6 differs from NAT66, which is stateful. Palo Alto Networks supports NPTv6 RFC 6296 prefix translation; it does not support NAT66.

With the limited addresses in the IPv4 space, NAT was required to translate private, non-routable IPv4 addresses to one or more globally-routable IPv4 addresses.

For organizations using IPv6 addressing, there is no need to translate IPv6 addresses to IPv6 addresses due to the abundance of IPv6 addresses. However, there are Reasons to Use NPTv6 to translate IPv6 prefixes at the firewall.

NPTv6 translates the prefix portion of an IPv6 address but not the host portion or the application port numbers. The host portion is simply copied, and therefore remains the same on either side of the firewall. The host portion also remains visible within the packet header.

- NPTv6 Does Not Provide Security
- Platform Support for NPTv6
- Unique Local Addresses
- Reasons to Use NPTv6

## NPTv6 Does Not Provide Security

It is important to understand that NPTv6 does not provide security. In general, stateless network address translation does not provide any security; it provides an address translation function. NPTv6 does not hide or translate port numbers. You must set up firewall security policies correctly in each direction to ensure that traffic is controlled as you intended.

## Platform Support for NPTv6

NPTv6 is supported on the following platforms (NPTv6 with hardware lookup but packets go through the CPU): PA-7000 Series, PA-5000 Series, PA-4000 Series, PA-3050 firewall, and PA-2000 Series. Platforms supported with no ability to have hardware perform a session look-up: PA-3020 firewall, PA 500 firewall, PA-200 firewall, and VM-Series.

## Unique Local Addresses

RFC 4193, Unique Local IPv6 Unicast Addresses, defines unique local addresses (ULAs), which are IPv6 unicast addresses. They can be considered IPv6 equivalents of the private IPv4 addresses identified in RFC 1918, Address Allocation for Private Internets, which cannot be routed globally.

 © Palo Alto Networks, Inc.

A ULA is globally unique, but not expected to be globally routable. It is intended for local communications and to be routable in a limited area such as a site or among a small number of sites. Palo Alto Networks does not recommend that you assign ULAs, but a firewall configured with NPTv6 will translate prefixes sent to it, including ULAs.

## Reasons to Use NPTv6

Although there is no shortage of public, globally routable IPv6 addresses, there are reasons you might want to translate IPv6 addresses. NPTv6:

- **Prevents asymmetrical routing**—Asymmetric routing can occur if a Provider Independent address space (/48, for example) is advertised by multiple data centers to the global Internet. By using NPTv6, you can advertise more specific routes from regional firewalls, and the return traffic will arrive at the same firewall where the source IP address was translated by the translator.
- **Provides address independence**—You need not change the IPv6 prefixes used inside your local network if the global prefixes are changed (for example, by an ISP or as a result of merging organizations). Conversely, you can change the inside addresses at will without disrupting the addresses that are used to access services in the private network from the Internet. In either case, you update a NAT rule rather than reassign network addresses.
- **Translates ULAs for routing**—You can have Unique Local Addresses assigned within your private network, and have the firewall translate them to globally routable addresses. Thus, you have the convenience of private addressing and the functionality of translated, routable addresses.
- **Reduces exposure to IPv6 prefixes**—IPv6 prefixes are less exposed than if you didn't translate network prefixes, however, NPTv6 is not a security measure. The interface identifier portion of each IPv6 address is not translated; it remains the same on each side of the firewall and visible to anyone who can see the packet header. Additionally, the prefixes are not secure; they can be determined by others.

© Palo Alto Networks, Inc.

## How NPTv6 Works

When you configure a policy for NPTv6, the Palo Alto Networks firewall performs a static, one-to-one IPv6 translation in both directions. The translation is based on the algorithm described in RFC 6296.

In one use case, the firewall performing NPTv6 is located between an internal network and an external network (such as the Internet) that uses globally routable prefixes. When datagrams are going in the outbound direction, the internal source prefix is replaced with the external prefix; this is known as source translation.

In another use case, when datagrams are going in the inbound direction, the destination prefix is replaced with the internal prefix (known as destination translation). The figure below illustrates destination translation and a characteristic of NPTv6: only the prefix portion of an IPv6 address is translated. The host portion of the address is not translated and remains the same on either side of the firewall. In the figure below, the host identifier is 111::55 on both sides of the firewall.

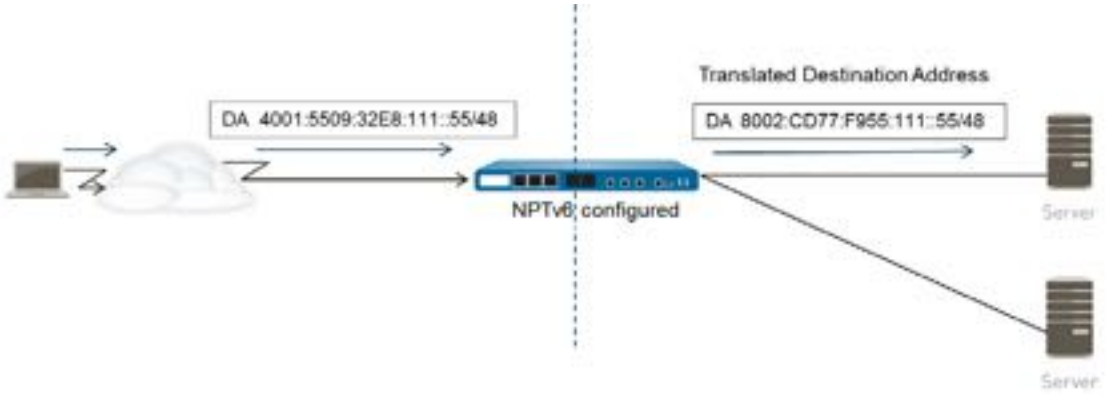


It is important to understand that NPTv6 does not provide security. While you are planning your NPTv6 NAT policies, remember also to configure security policies in each direction.

A NAT or NPTv6 policy rule cannot have both the Source Address and the Translated Address set to Any.

In an environment where you want IPv6 prefix translation, three firewall features work together: NPTv6 NAT policies, security policies, and NDP Proxy.

The firewall does not translate the following:

- Addresses that the firewall has in its Neighbor Discovery (ND) cache.
- The subnet 0xFFFF (in accordance with RFC 6296, Appendix B).
- IP multicast addresses.
- IPv6 addresses with a prefix length of /31 or shorter.
- Link-local addresses. If the firewall is operating in virtual wire mode, there are no IP addresses to translate, and the firewall does not translate link-local addresses.
- Addresses for TCP sessions that authenticate peers using the TCP Authentication Option (RFC 5925).

When using NPTv6, performance for fast path traffic is impacted because NPTv6 is performed in the slow path.

 © Palo Alto Networks, Inc.

NPTv6 will work with IPSec IPv6 only if the firewall is originating and terminating the tunnel. Transit IPSec traffic would fail because the source and/or destination IPv6 address would be modified. A NAT traversal technique that encapsulates the packet would allow IPSec IPv6 to work with NPTv6.


- Checksum-Neutral Mapping
- Bi-Directional Translation
- NPTv6 Applied to a Specific Service


## Checksum-Neutral Mapping

The NPTv6 mapping translations that the firewall performs are checksum-neutral, meaning that "... they result in IP headers that will generate the same IPv6 pseudo-header checksum when the checksum is calculated using the standard Internet checksum algorithm [RFC 1071]." See RFC 6296, Section 2.6, for more information about checksum-neutral mapping.

If you are using NPTv6 to perform destination NAT, you can provide the internal IPv6 address and the external prefix/prefix length of the firewall interface in the syntax of the `test nptv6` CLI command. The CLI responds with the checksum-neutral, public IPv6 address to use in your NPTv6 configuration to reach that destination.

## Bi-Directional Translation

When you Create an NPTv6 Policy, the **Bi-directional** check box in the **Translated Packet** tab provides a convenient way for you to have the firewall create a corresponding NAT or NPTv6 translation in the opposite direction of the translation you configured. By default, **Bi-directional** translation is disabled.

If you enable **Bi-directional** translation, it is very important to make sure you have security policies in place to control the traffic in both directions. Without such policies, the **Bi-directional** feature will allow packets to be automatically translated in both directions, which you might not want.

## NPTv6 Applied to a Specific Service

The Palo Alto Networks implementation of NPTv6 offers the ability to filter packets to limit which packets are subject to translation. Keep in mind that NPTv6 does not perform port translation. There is no concept of Dynamic IP and Port (DIPP) translation because NPTv6 translates IPv6 prefixes only. However, you can specify that only packets for a certain service port undergo NPTv6 translation. To do so, Create an NPTv6 Policy that specifies a **Service** in the Original Packet.

## NDP Proxy

Neighbor Discovery Protocol (NDP) for IPv6 performs functions similar to those provided by Address Resolution Protocol (ARP) for IPv4. RFC 4861 defines Neighbor Discovery for IP version 6 (IPv6). Hosts, routers, and firewalls use NDP to determine the link-layer addresses of neighbors on connected links, to keep track of which neighbors are reachable, and to update neighbors' link-layer addresses that have changed. Peers advertise their own MAC address and IPv6 address, and they also solicit addresses from peers.

NDP also supports the concept of *proxy*, when a node has a neighboring device that is able to forward packets on behalf of the node. The device (firewall) performs the role of NDP Proxy.

Palo Alto Networks firewalls support NDP and NDP Proxy on their interfaces. When you configure the firewall to act as an NDP Proxy for addresses, it allows the firewall to send Neighbor Discovery (ND) advertisements and respond to ND solicitations from peers that are asking for MAC addresses of IPv6 prefixes assigned to devices behind the firewall. You can also configure addresses for which the firewall will not respond to proxy requests (negated addresses).

In fact, NDP is enabled by default, and you need to configure NDP Proxy when you configure NPTv6, for the following reasons:

- The stateless nature of NPTv6 requires a way to instruct the firewall to respond to ND packets sent to specified NDP Proxy addresses, and to not respond to negated NDP Proxy addresses.

It is recommended that you negate your neighbors' addresses in the NDP Proxy configuration, because NDP Proxy indicates the firewall will reach those addresses behind the firewall, but the neighbors are not behind the firewall.

- NDP causes the firewall to save the MAC addresses and IPv6 addresses of neighbors in its ND cache. (Refer to the figure in NPTv6 and NDP Proxy Example.) The firewall does not perform NPTv6 translation for addresses that it finds in its ND cache because doing so could introduce a conflict. If the host portion of an address in the cache happens to overlap with the host portion of a neighbor's address, and the prefix in the cache is translated to the same prefix as that of the neighbor (because the egress interface on the firewall belongs to the same subnet as the neighbor), then you would have a translated address that is exactly the same as the legitimate IPv6 address of the neighbor, and a conflict occurs. (If an attempt to perform NPTv6 translation occurs on an address in the ND cache, an informational syslog message logs the event: `NPTv6 Translation Failed.`)

When an interface with NDP Proxy enabled receives an ND solicitation requesting a MAC address for an IPv6 address, the following sequence occurs:

1. The firewall searches the ND cache to ensure the IPv6 address from the solicitation is not there. If the address is there, the firewall ignores the ND solicitation.
2. If the source IPv6 address is 0, that means the packet is a Duplicate Address Detection packet, and the firewall ignores the ND solicitation.
3. The firewall does a Longest Prefix Match search of the NDP Proxy addresses and finds the best match to the address in the solicitation. If the Negate field for the match is checked (in the NDP Proxy list), the firewall drops the ND solicitation.
4. Only if the Longest Prefix Match search matches, and that matched address is not negated, will the NDP Proxy respond to the ND solicitation. The firewall responds with an ND packet, providing its own MAC address as the MAC address of the next hop toward the queried destination.

 © Palo Alto Networks, Inc.

In order to successfully support NDP, the firewall does not perform NDP Proxy for the following:

- Duplicate Address Detection (DAD).
- Addresses in the ND cache (because such addresses do not belong to the firewall; they belong to discovered neighbors).

© Palo Alto Networks, Inc.

## NPTv6 and NDP Proxy Example

The following figure and text illustrate how NPTv6 and NDP Proxy function together.

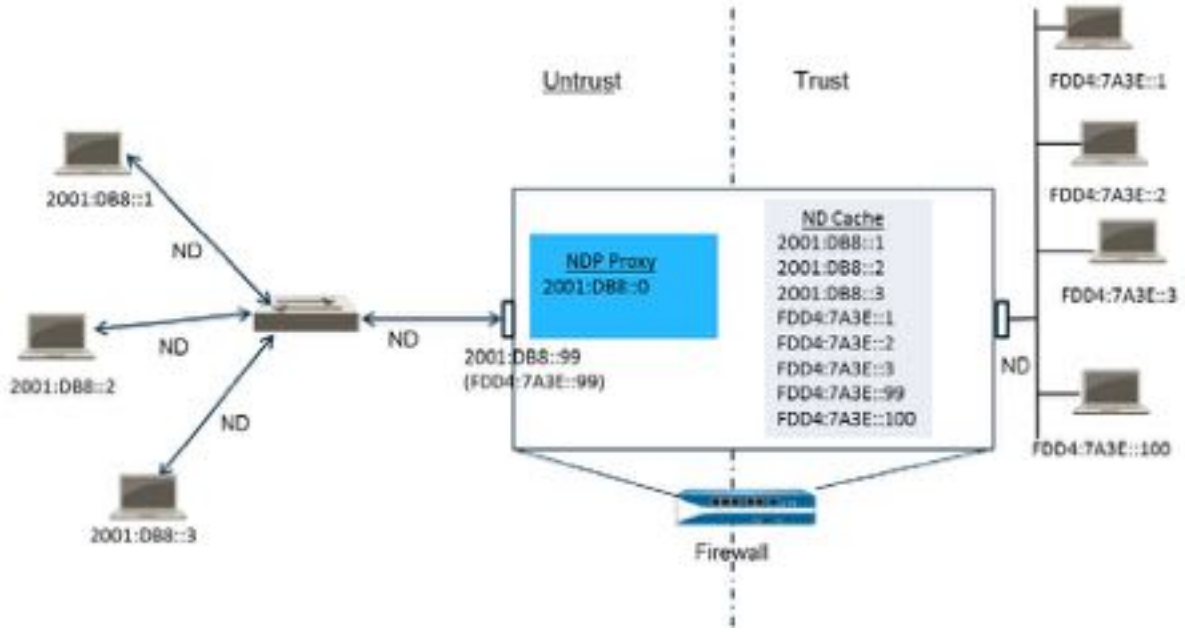


### The ND Cache in NPTv6 Example

In the above example, multiple peers connect to the firewall though a switch, with ND occurring between the peers and the switch, between the switch and the firewall, and between the firewall and the devices on the trust side.

As the firewall learns of peers, it saves their addresses to its ND cache. Trusted peers FDDA:7A3E::1, FDDA:7A3E::2, and FDDA:7A3E::3 are connected to the firewall on the trust side. FDDA:7A3E::99 is the untranslated address of the firewall itself; its public-facing address is 2001:DB8::99. The addresses of the peers on the untrust side have been discovered and appear in the ND cache: 2001:DB8::1, 2001:DB8::2, and 2001:DB8::3.

### The NDP Proxy in NPTv6 Example

In our scenario, we want the firewall to act as NDP Proxy for the prefixes on devices behind the firewall. When the firewall is NDP Proxy for a specified set of addresses/ranges/prefixes, and it sees an address from this range in an ND solicitation or advertisement, the firewall will respond as long as a device with that specific address doesn't respond first, the address is not negated in the NDP proxy configuration, and the address is not in the ND cache. The firewall does the prefix translation (described below) and sends the packet to the trust side, where that address might or might not be assigned to a device.

 © Palo Alto Networks, Inc.

In this example, the ND Proxy table contains the network address 2001:DB8::0. When the interface sees an ND for 2001:DB8::100, no other devices on the L2 switch claim the packet, so the proxy range causes the firewall to claim it, and after translation to FDD4:7A3E::100, the firewall sends it out to the trust side.

## The NPTv6 Translation in NPTv6 Example

In this example, the **Original Packet** is configured with a **Source Address** of FDD4:7A3E::0 and a **Destination** of **Any**. The **Translated Packet** is configured with the **Translated Address** of 2001:DB8::0.

Therefore, outgoing packets with a source of FDD4:7A3E::0 are translated to 2001:DB8::0. Incoming packets with a destination prefix in the network 2001:DB8::0 are translated to FDD4:7A3E::0.

## Neighbors in the ND Cache are Not Translated

In our example, there are hosts behind the firewall with host identifiers :1, :2, and :3. If the prefixes of those hosts are translated to a prefix that exists beyond the firewall, and if those devices also have host identifiers :1, :2, and :3, because the host identifier portion of the address remains unchanged, the resulting translated address would belong to the existing device, and an addressing conflict would result. In order to avoid a conflict with overlapping host identifiers, NPTv6 does not translate addresses that it finds it its ND cache.

## Create an NPTv6 Policy

Perform this task when you want to configure a NAT NPTv6 policy to translate one IPv6 prefix to another IPv6 prefix. The prerequisites for this task are:

- Enable IPv6. Select **Device > Setup > Session**. Click **Edit** and select **IPv6 Firewalling**.
- Configure a Layer 3 Ethernet interface with a valid IPv6 address and with IPv6 enabled. Select **Network > Interfaces > Ethernet**, select an interface, and on the **IPv6** tab, select **Enable IPv6 on the interface**.
- Create network security policies, because NPTv6 does not provide security.
- Decide whether you want source translation, destination translation, or both.
- Identify the zones to which you want to apply the NPTv6 policy.
- Identify your original and translated IPv6 prefixes.

| **Configure an NPTv6 Policy** | | |
|---|---|---|
| Step 1 | Create a new NPTv6 policy. | 1. Select **Policies** > **NAT** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the NPTv6 policy rule.<br>3. (Optional) Enter a **Description** and **Tag**.<br>4. For **NAT Type**, select **NPTv6**. |
| Step 2 | Specify the match criteria for incoming packets; packets that match all of the criteria are subject to the NPTv6 translation.<br><br>Zones are required for both types of translation. | 1. On the **Original Packet** tab, for **Source Zone**, leave **Any** or click **Add** to enter the source zone to which the policy applies.<br>2. Enter the **Destination Zone** to which the policy applies.<br>3. (Optional) Select a **Destination Interface**.<br>4. (Optional) Select a **Service** to restrict what type of packets are translated.<br>5. If you are doing source translation, enter a **Source Address** or select **Any**. The address could be an address object. The following constraints apply to **Source Address** and **Destination Address**:<br>• Prefixes of **Source Address** and **Destination Address** for the **Original Packet** and **Translated Packet** must be in the format xxxx:xxxx::/yy, although leading zeros in the prefix can be dropped.<br>• The IPv6 address cannot have an interface identifier (host) portion defined.<br>• The range of supported prefix lengths is /32 to /64.<br>• The **Source Address** and **Destination Address** cannot both be set to **Any**.<br>6. If you are doing source translation, you can optionally enter a **Destination Address**. If you are doing destination translation, the **Destination Address** is required. See the constraints listed in the prior step. |


© Palo Alto Networks, Inc.

| Configure an NPTv6 Policy (Continued) | | |
|---|---|---|
| Step 3 | Specify the translated packet. | 1. On the **Translated Packet** tab, if you want to do source translation, in the Source Address Translation section, for **Translation Type**, select **Static IP**. If you do not want to do source translation, select **None**.<br>2. If you chose **Static IP**, the **Translated Address** field appears. Enter the translated IPv6 prefix or address object. See the constraints listed in Step 5.<br>It is a best practice to configure your **Translated Address** to be the prefix of the untrust interface address of your firewall. For example, if your untrust interface has the address 2001:1a:1b:1::99/64, make your **Translated Address** 2001:1a:1b:1::0/64.<br>3. (Optional) Select **Bi-directional** if you want the firewall to create a corresponding NPTv6 translation in the opposite direction of the translation you configure.<br>If you enable **Bi-directional** translation, it is very important to make sure you have security policies in place to control the traffic in both directions. Without such policies, the **Bi-directional** feature will allow packets to be automatically translated in both directions, which you might not want.<br>4. If you want to do destination translation, select **Destination Address Translation**. In the **Translated Address** field, choose an address object from the drop-down or enter your internal destination address.<br>5. Click **OK**. |
| Step 4 | Configure NDP Proxy.<br>When you configure the firewall to act as an NDP Proxy for addresses, it allows the firewall to send Neighbor Discovery (ND) advertisements and respond to ND solicitations from peers that are asking for MAC addresses of IPv6 prefixes assigned to devices behind the firewall. | 1. Select **Network > Interfaces > Ethernet** and select an interface.<br>2. On the **Advanced** > **NDP Proxy** tab, select **Enable NDP Proxy** and click **Add**.<br>3. Enter the **IP Address(es)** for which NDP Proxy is enabled. It can be an address, a range of addresses, or a prefix and prefix length. The order of IP addresses does not matter. These addresses are ideally the same as the Translated Addresses that you configured in an NPTv6 policy.<br>If the address is a subnet, the NDP Proxy will respond to all addresses in the subnet, so you should list the neighbors in that subnet with **Negate** selected, as described in the next step.<br>4. (Optional) Enter one or more addresses for which you do not want NDP Proxy enabled, and select **Negate**. For example, from an IP address range or prefix range configured in the prior step, you could negate a smaller subset of addresses. It is recommended that you negate the addresses of the neighbors of the firewall. |
| Step 5 | Save the configuration. | Click **OK** and **Commit**. |

# LACP

The firewall can now use Link Aggregation Control Protocol (LACP) to detect the physical interfaces between itself and a connected device (peer) and manage those interfaces as a single virtual interface (aggregate group). An aggregate group increases the bandwidth between peers. Enabling LACP provides redundancy within the group: the protocol automatically detects interface failures and performs failover to standby interfaces. Without LACP, you must manually identify interface failures occurring at layers above the physical or occurring between peers that are not directly connected.

The following Palo Alto Networks firewalls support LACP: PA-500, PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7000 Series. The firewalls support LACP for HA3 (only on the PA-500, PA-3000 Series, PA-4000 Series, and PA-5000 Series), Layer 2, and Layer 3 interfaces.

The following topics describe LACP and how to configure it for aggregate groups:

- LACP Settings
- Configure LACP

## LACP Settings

To implement LACP, you must configure an aggregate group on each LACP peer that will use the protocol. Before you Configure LACP, determine the optimal settings for each peer, as described in the following topics:

- Mode
- Transmission Rate
- Fast Failover
- Port Priority and System Priority

### Mode

In LACP active mode, a device actively searches the network for peers and queries their status (available or unresponsive). In passive mode, a device only responds to queries from an active peer. Between any two peers, it is recommended that one be active and the other passive. LACP cannot function if both peers are passive.

### Transmission Rate

You can configure a device to detect the state (available or unresponsive) of peers and individual interfaces at fast (every second) or slow (every 30 seconds) intervals. When the device does not receive an LACP update from the peer within a period that is three times the transmission rate (3 seconds or 90 seconds), it flags the peer as unresponsive. Therefore, set the transmission rate according to how much LACP processing your network can support and how quickly a device should detect and resolve interface failures.

 © Palo Alto Networks, Inc.

For example, in a network where peers have many redundant links, a single interface failure would not disrupt traffic. In this case, a slow rate would suffice for detecting failures and would avoid the extra processing associated with the fast rate. If the network can support the extra processing and requires high availability (for example, datacenters supporting critical business operations), a fast transmission rate would enable the firewall to substitute failed interfaces quicker.

If devices have different transmission rates, each uses the rate of its peer.

## Fast Failover

When an interface goes down, it fails over to a standby interface. The IEEE 802.1ax standard that defines LACP specifies a failover process that takes at least three seconds. For networks that require quicker failover, the firewall provides an option that performs fast failover within one second. This option is recommended for deployments in which critical data might be lost during the standard failover interval. For example, if the amount of traffic between peers is close to the maximum that the active interfaces can support and one interface fails, the risk of data being lost while a standby interface becomes active is significantly less for fast failover than for standard failover.

Because both peers perform failover when an interface goes down, it is recommended that you configure fast failover on both so that they complete the process within the same time frame.

## Port Priority and System Priority

When you configure an aggregate group, the **Max Ports** parameter determines how many interfaces can be active at any given time. If the number of interfaces you assign to the group exceeds the **Max Ports**, the remaining interfaces are in standby mode. If an active interface fails, a standby interface becomes active. LACP uses the **LACP Port Priority** assigned to each interface to determine which interfaces are initially active and to determine the order in which standby interfaces become active upon failover. (You set port priorities when configuring individual interfaces, not when configuring the aggregate group.)

It is recommended that you assign a unique priority number to each interface in an aggregate group. However, if multiple interfaces have the same priority number, the interface with the lowest port number has the highest priority.

To enable LACP between two peers, each requires an aggregate group configuration. If the port priority values for the member interfaces differ in each peer, the values of the peer with the higher **System Priority** will override the other peer in determining active and standby interfaces.

For port and system priorities, the priority level is the inverse of its numerical value: a priority of 1 designates the highest priority level; a value of 65535 designates the lowest priority level. The default is 32768.

The following figure shows an aggregate group instance configured on a firewall and on a switch. The group has four interfaces and the **Max Ports** parameter is set to three. That means three interfaces are active and one is standby. The firewall administrator assigned port priorities 12, 64, 128, and 258 to the interfaces. The switch administrator assigned port priorities 22, 54, 158, and 228 to the interfaces. LACP uses the port priorities of the firewall instance because the **System Priority** of the firewall is higher than that of the switch.

**Figure: LACP System Priority**

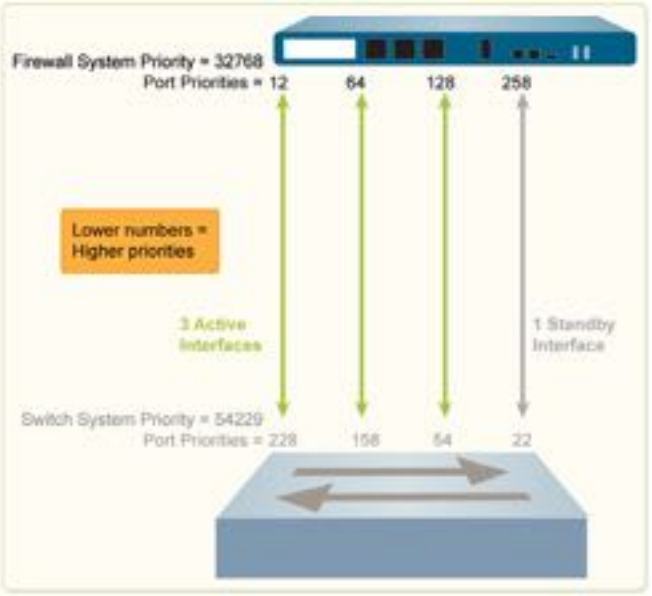


It is a best practice to assign a unique system priority to each peer. However, in some cases, peers might have the same value. This could happen if a different administrator configured the aggregate group on each peer. In such cases, the peer for which the *system ID* is a lower numerical value will override the other peer with respect to port priorities. LACP automatically derives the system ID from the system priority and system MAC address. A numerically lower MAC address produces a numerically lower system ID.

The following figure shows the same peers as in Figure: LACP System Priority but with identical system priorities. In this case, LACP uses the system IDs to determine port prioritization: the switch overrides the firewall.

**Figure: LACP System ID**

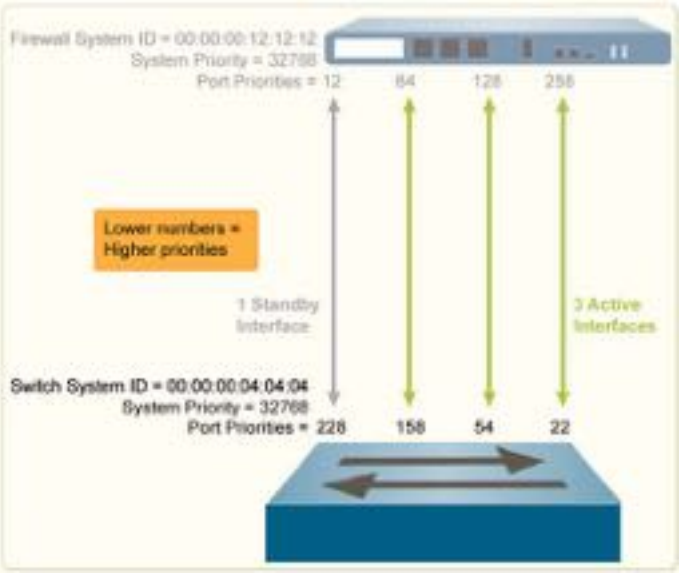


 © Palo Alto Networks, Inc.

Firewalls in a high availability (HA) pair have the same system priority value. However, in an active/passive deployment, the system ID for each can be the same or different, depending on whether you assign the same MAC address. (Firewalls in an active/active deployment require unique MAC addresses so PAN-OS automatically assigns them.) When the LACP peers (also in HA mode) are virtualized (appearing to the network as a single device), selecting the **Same System MAC Address for Active-Passive HA** option for the firewalls is a best practice to minimize latency during failover. When the LACP peers are not virtualized, using the unique MAC address of each firewall is the best practice to minimize failover latency. In the latter case, if the HA firewall pair and the LACP peer devices all have the same system priority but each firewall in the HA pair has a unique system ID, one firewall might have a lower system ID than the LACP peers while the other firewall has a higher system ID than the LACP peers. In this case, when failover occurs on the firewalls, port prioritization switches between the LACP peers and the firewall that becomes active.

## Configure LACP

Before starting this procedure:

- Determine which physical interfaces connect the LACP peers. This procedure assumes the cabling is complete.
- Determine the optimal LACP Settings for the peers.

This procedure only covers the configuration steps for an LACP peer that is a Palo Alto Networks firewall. Other devices have proprietary methods for configuring LACP.

Perform the following steps to configure LACP on a firewall. In a high availability deployment, configure the primary (for active/active) or active (for active/passive) firewall; the secondary or passive firewall automatically synchronizes with it.

| Configure LACP | |
|---|---|
| Step 1 Add an aggregate group with LACP enabled. | 1. Select **Network > Interfaces > Ethernet** and click **Add Aggregate Group**.<br>2. In the field adjacent to the read-only **Interface Name**, enter a number (1-8) to identify the group.<br>3. For the **Interface Type**, select **HA**, **Layer 2**, or **Layer 3**.<br>Only select **HA** if the interface is an HA3 link between two firewalls in an active/active deployment.<br>4. On the **LACP** tab, select **Enable LACP**.<br>5. For the LACP **Mode**, select whether the firewall is **Passive** (default) or **Active**.<br>6. For the **Transmission Rate**, select **Fast** or **Slow**.<br>7. If desired, select **Fast Failover** (under one second) for when interfaces go down. Otherwise, the standard failover (at least three seconds) applies.<br>8. Enter a **System Priority** number (1-65535, default 32768) that determines whether the firewall or peer overrides the other with respect to port priorities. Note that the lower the number, the higher the priority.<br>9. For the **Max Ports**, enter the number of interfaces (1-8) that are active. The value cannot exceed the number of interfaces you assign to the group. If the number of assigned interfaces exceeds the number of active interfaces, the remaining interfaces will be in standby mode.<br>10. By default, each firewall in an HA pair has a unique MAC address. A **Same System MAC Address for Active-Passive HA** configuration is recommended only when the LACP peers are virtualized (appearing to the network as a single device). In such cases, select the check box and select the system-generated **MAC Address**, or enter your own. You must verify the address is globally unique.<br>If the firewalls are not in active/passive HA mode, PAN-OS ignores your selections for these fields. Firewalls in an active/active deployment require unique MAC addresses so PAN-OS automatically assigns them. If the **Interface Type** is **HA**, these fields do not appear.<br>11. Click **OK**. |

 © Palo Alto Networks, Inc.

| Configure LACP (Continued) | | |
|---|---|---|
| Step 2 | Assign interfaces to the aggregate group. | Perform the following steps for each physical interface (1-8) that will belong to the aggregate group.<br>During normal operations, assign more than one interface. You would only assign a single interface if a troubleshooting procedure requires it.<br>1. Select **Network > Interfaces** and select the interface name. It must be the same type (HA, Layer 2, or Layer 3) as the aggregate group.<br>2. Change the **Interface Type** to **Aggregate Ethernet**.<br>3. Select the **Aggregate Group** you just defined.<br>4. Select the **Link Speed**, **Link Duplex**, and **Link State**. It is a best practice to set the same link speed and duplex values for every interface in the group. For non-matching values, the commit operation displays a warning and PAN-OS defaults to the higher speed and full duplex.<br>5. Enter an **LACP Port Priority** (1-65535, default 32768). If the number of interfaces you assign exceeds the **Max Ports** value you configured for the group, the port priorities determine which interfaces are active or standby. The lower the numeric value, the higher the priority.<br>6. Click **OK** and **Commit**. |
| Step 3 | Verify the status of the aggregate group. | 1. Select **Network > Interfaces > Ethernet**.<br>2. Verify that the Features column displays an LACP enabled icon for the aggregate group.<br>3. Verify that the Link State column displays a green icon for the aggregate group, indicating that all member interfaces are up. If the icon is yellow, at least one member is down but not all. If the icon is red, all members are down.<br>4. If the Link State icon for the aggregate group is yellow or red, select **Monitor > Logs > System** and review the logs for the LACP subtype to investigate the causes of the interface failures. |
| Step 4 | (Optional) If you selected **HA** as the **Interface Type** for the aggregate group, enable forwarding of packets over the HA3 link. | 1. Select **Device > High Availability > Active/Active Config** and edit the Packet Forwarding section.<br>2. For the **HA3 Interface**, select the aggregate group you configured.<br>3. Click **OK** and **Commit**. |

# ECMP

Equal Cost Multiple Path (ECMP) processing is a networking feature that enables the firewall to use up to four equal-cost routes to the same destination. Without this feature, if there are multiple equal-cost routes to the same destination, the virtual router chooses one of those routes from the routing table and adds it to its forwarding table; it will not use any of the other routes unless there is an outage in the chosen route.

Enabling ECMP functionality on a virtual router allows the firewall to have up to four equal-cost paths to a destination in its forwarding table, allowing the firewall to:

- Load balance flows (sessions) to the same destination over multiple equal-cost links.
- Efficiently use all available bandwidth on links to the same destination rather than leave some links unused.
- Dynamically shift traffic to another ECMP member to the same destination if a link fails, rather than having to wait for the routing protocol or RIB table to elect an alternative path/route. This can help reduce downtime when links fail.

The following sections describe ECMP and how to configure it.


- ECMP Load-Balancing Algorithms
- ECMP Platform, Interface, and IP Routing Support
- HA Active/Active Failover Behavior with ECMP
- Configure ECMP on a Virtual Router
- Enable ECMP for Multiple BGP Autonomous Systems
- Verify ECMP


 © Palo Alto Networks, Inc.

## ECMP Load-Balancing Algorithms

Let's suppose the Routing Information Base (RIB) of the firewall has multiple equal-cost paths to a single destination. The maximum number of equal-cost paths defaults to 2. ECMP chooses the best two equal-cost paths from the RIB to copy to the Forwarding Information Base (FIB). ECMP then determines, based on the load-balancing method, which of the two paths in the FIB that the firewall will use for the destination during this session.

ECMP load balancing is done at the session level, not at the packet level—the start of a new session is when the firewall (ECMP) chooses an equal-cost path. The equal-cost paths to a single destination are considered ECMP path members or ECMP group members. ECMP determines which one of the multiple paths to a destination in the FIB to use for an ECMP flow, based on which load-balancing algorithm you set. A virtual router can use only one load-balancing algorithm.

Enabling, disabling, or changing ECMP on an existing virtual router causes the system to restart the virtual router, which might cause existing sessions to be terminated.

The four algorithm choices emphasize different priorities, as follows:

- **Hash-based algorithms prioritize session stickiness**—The **IP Modulo** and **IP Hash** algorithms use hashes based on information in the packet header, such as source and destination address. Because the header of each flow in a given session contains the same source and destination information, these options prioritize session *stickiness*. If you choose the **IP Hash** algorithm, you can optionally set a **Hash Seed** value to further randomize load balancing if you have a large number of sessions to the same destination and they're not being distributed evenly over the ECMP links.

- **Balanced algorithm prioritizes load balancing**—The **Balanced Round Robin** algorithm distributes incoming sessions equally across the links, favoring load balancing over session stickiness. (Round robin indicates a sequence in which the least recently chosen item is chosen.) In addition, if new routes are added or removed from an ECMP group (for example if a path in the group goes down), the virtual router will re-balance the sessions across links in the group. Additionally, if the flows in a session have to switch routes due to an outage, when the original route associated with the session becomes available again, the flows in the session will revert to the original route when the virtual router once again re-balances the load.

- **Weighted algorithm prioritizes link capacity and/or speed**—As an extension to the ECMP protocol standard, the Palo Alto Networks implementation provides for a **Weighted Round Robin** load-balancing option that takes into account differing link capacities and speeds on the egress interfaces of the firewall. With this option, you can assign **ECMP Weights** (range is 1-255; default is 100) to the interfaces based on link performance using factors such as link capacity, speed, and latency to ensure that loads are balanced to fully leverage the available links.

  For example, suppose the firewall has redundant links to an ISP: ethernet1/1 (100 Mbps) and ethernet1/8 (200 Mbps). Although these are equal-cost paths, the link via ethernet1/8 provides greater bandwidth and therefore can handle a greater load than the ethernet1/1 link. Therefore, to ensure that the load-balancing functionality takes into account link capacity and speed, you might assign ethernet1/8 a weight of 200 and ethernet1/1 a weight of 100. The 2:1 weight ratio causes the virtual router to send twice as many sessions to ethernet1/8 as it sends to ethernet1/1. However, because the ECMP protocol is inherently session-based, when using the **Weighted Round Robin** algorithm, the firewall will be able to load balance across the ECMP links only on a best-effort basis.

© Palo Alto Networks, Inc.

Assign lower-speed or lower-capacity links with a lower weight. Assign higher-speed or higher-capacity links with a higher weight. In this manner, the firewall can distribute sessions based on these ratios, rather than overdrive a low-capacity link that is one of the equal-cost paths.

Keep in mind that ECMP weights are assigned to interfaces to determine load balancing (to influence which *equal-cost* path is chosen), not for route selection (a route choice from routes that could have different costs).

 © Palo Alto Networks, Inc.

## ECMP Platform, Interface, and IP Routing Support

ECMP is supported on all Palo Alto Networks firewall platforms, with hardware forwarding support on the PA-7000 Series, PA-5000 Series, PA-3060 firewalls, and PA-3050 firewalls. PA-3020 firewalls, PA-500 firewalls, PA-200 firewalls, and VM-Series firewalls support ECMP through software only. Performance is affected for sessions that cannot be hardware offloaded.

ECMP is supported on Layer 3, Layer 3 subinterface, VLAN, tunnel, and Aggregated Ethernet interfaces.

ECMP can be configured for static routes and any of the dynamic routing protocols the firewall supports.

ECMP affects the route table capacity because the capacity is based on the number of paths, so an ECMP route with four paths will consume four entries of route table capacity. ECMP implementation might slightly decrease the route table capacity because more memory is being used by session-based tags to map traffic flows to particular interfaces.

ECMP has the following restrictions:

- PA-2000 Series firewalls and PA-4000 Series firewalls with ECMP enabled might not be able to offload sessions to hardware for forwarding. Packets matching ECMP routes will be sent to software, while packets matching non-ECMP routes can still be forwarded by hardware.
- For the PA-4000 Series firewalls, packets to be forwarded by ECMP routes will be sent to software for route lookup and forwarding, even though the session is in offloaded state.
- Virtual router-to-virtual router routing using static routes does not support ECMP.

© Palo Alto Networks, Inc.

## HA Active/Active Failover Behavior with ECMP

In the event of a failover in a high availability (HA) active/active configuration, the sessions on the failed peer are transferred to the HA peer. When this happens, the new active primary device tries to use the same egress interface that the former failed active/active device was using. If such interface is found among the ECMP paths, the transferred sessions will take the same egress interface and path. This behavior takes place regardless of the ECMP algorithm in use; using the same interface is desirable.

Only if no ECMP path matches the original egress interface will the active-primary select a new ECMP path.

In an HA active/active configuration, if you did not configure the same interfaces on both peers, when a failover occurs, the active-primary will select the next best path from the FIB table. In this scenario, the existing sessions might not be distributed according to the ECMP algorithm.

 © Palo Alto Networks, Inc.

## Configure ECMP on a Virtual Router

Use the following procedure to enable ECMP on a virtual router. The prerequisites are to:

- ❒ Specify the interfaces that belong to a virtual router (**Network > Virtual Routers > Router Settings > General**).
- ❒ Specify the IP routing protocol.

Enabling, disabling, or changing ECMP for an existing virtual router causes the system to restart the virtual router, which might cause sessions to be terminated.

| **Configure ECMP on a Virtual Router** | | |
|---|---|---|
| Step 1 | Enable ECMP for a virtual router. | 1. Select **Network > Virtual Routers** and select the virtual router on which to enable ECMP.<br>2. Select **Router Settings > ECMP** and select the **Enable** check box. |
| Step 2 | (Optional) Enable symmetric return of packets from server to client. | (Optional) Select the **Symmetric Return** check box to cause return packets to egress out the same interface on which the associated ingress packets arrived. That is, the firewall will use the ingress interface on which to send return packets, rather than use the ECMP interface. The **Symmetric Return** setting overrides load balancing. This behavior occurs only for traffic flows from the server to the client. |
| Step 3 | Specify the maximum number of equal-cost paths (to a destination network) that can be copied from the Routing Information Base (RIB) to the Forwarding Information Base (FIB). | For **Max Path** allowed, enter **2**, **3**, or **4**. Default: 2. |
| Step 4 | Select the load-balancing algorithm for the virtual router. For more information on load-balancing methods and how they differ, see ECMP Load-Balancing Algorithms. | For **Load Balance**, select one of the following options from the **Method** drop-down:<br>• **IP Modulo** (default)—Uses a hash of the source and destination IP addresses in the packet header to determine which ECMP route to use.<br>• **IP Hash**—Uses a hash of the source and destination IP addresses and optionally the source and destination port numbers in the packet header to determine which ECMP route to use. Specify options in Step 5 below.<br>• **Balanced Round Robin**—Uses round robin among the ECMP paths and re-balances paths when the number of paths changes.<br>• **Weighted Round Robin**—Uses round robin and a relative weight to select from among ECMP paths. Specify the weights in Step 6 below. |

| Configure ECMP on a Virtual Router (Continued) | |
|---|---|
| Step 5 (Optional) (**IP Hash** only) Configure IP Hash options. | If you selected **IP Hash** as the **Method**:<br>1. Check the **Use Source/Destination Ports** check box if you want to use source or destination port numbers in the **IP Hash** calculation.<br>2. Enter a **Hash Seed** value (an integer with a maximum of 9 digits). Specify a **Hash Seed** value to further randomize load balancing. Specifying a hash seed value is useful if you have a large number of sessions with the same tuple information. |
| Step 6 (**Weighted Round Robin** only) Define a weight for each interface in the ECMP group. | If you selected **Weighted Round Robin** as the **Method**, define a weight for each of the interfaces that are the egress points for traffic to be routed to the same destinations (that is, interfaces that are part of an ECMP group, such as the interfaces that provide redundant links to your ISP or interfaces to the core business applications on your corporate network).<br>The higher the weight, the more often that equal-cost path will be selected for a new session.<br>A higher speed link should be given a higher weight than a slower link so that more of the ECMP traffic goes over the faster link.<br>1. Create an ECMP group by clicking **Add** and selecting an **Interface** from the drop-down.<br>2. **Add** the other interfaces in the ECMP group.<br>3. Click on **Weight** and specify the relative weight for each interface (range is 1-255; default is 100). |
| Step 7 Save the configuration. | 1. Click **OK**.<br>2. Click **Yes** to restart the virtual router. Restarting the virtual router might cause existing sessions to be terminated.<br> <br>This message displays only if you are modifying an existing virtual router with ECMP. |
| Step 8 Save the configuration. | **Commit** the configuration. |

 © Palo Alto Networks, Inc.

## Enable ECMP for Multiple BGP Autonomous Systems

Perform the following task if you have BGP configured, and you want to enable ECMP over multiple autonomous systems. This task presumes that BGP is already configured. In the following figure, two ECMP paths to a destination go through two firewalls belonging to a single ISP in a single BGP autonomous system.

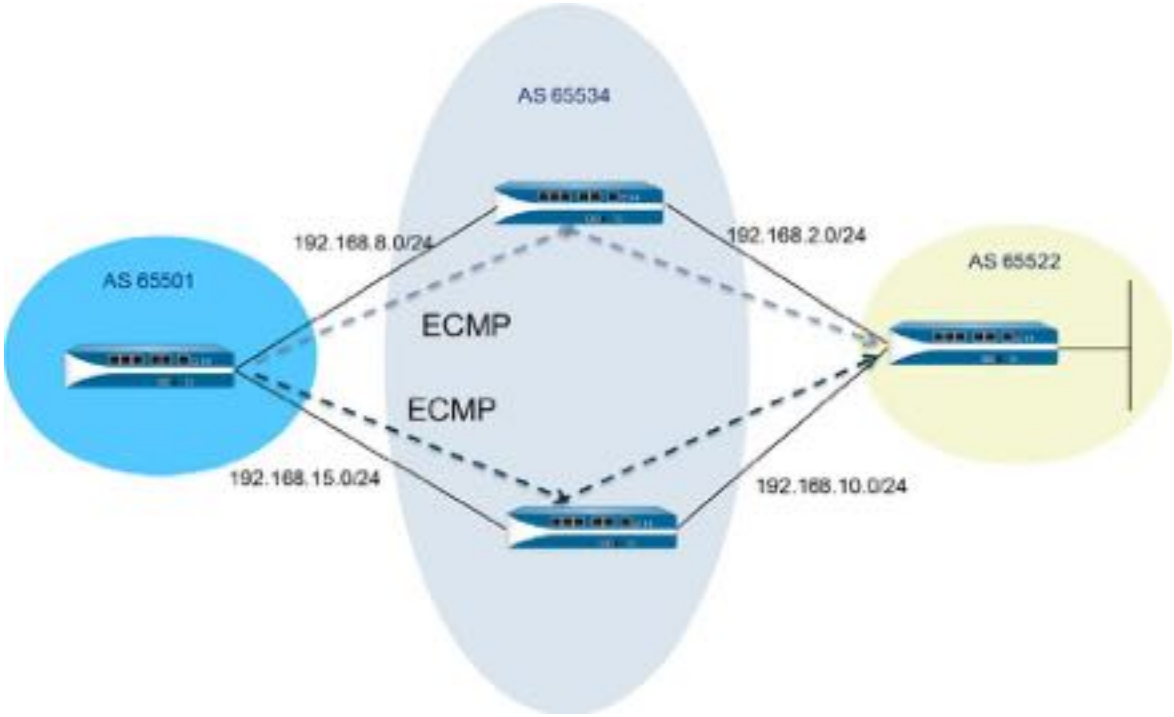


In the following figure, two ECMP paths to a destination go through two firewalls belonging to two different ISPs in different BGP autonomous systems.

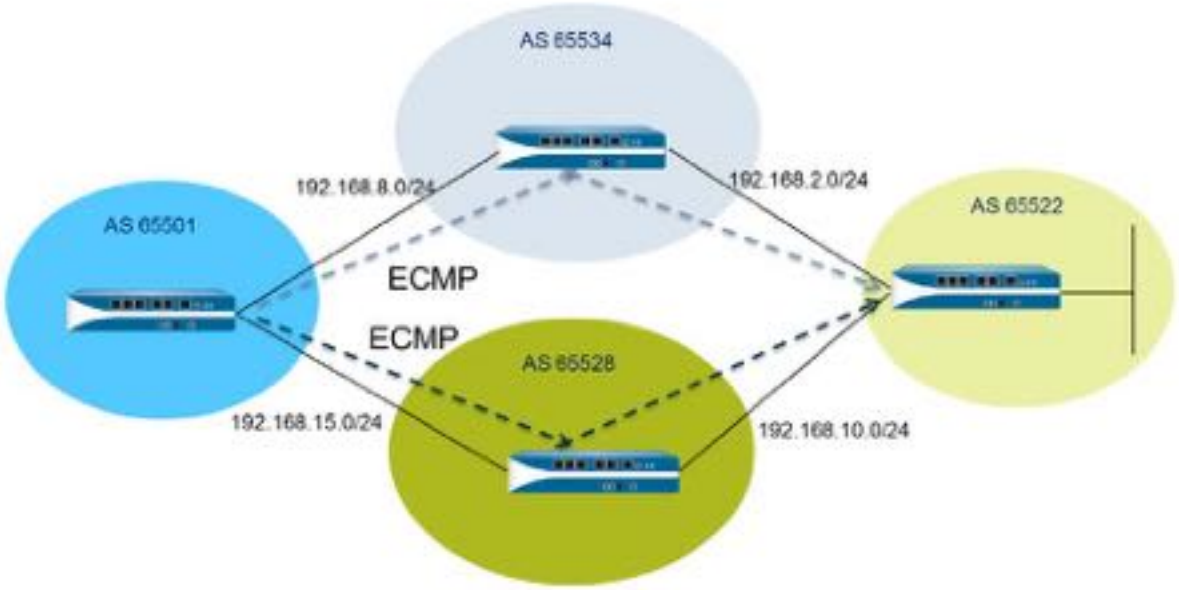


| Enable ECMP for BGP Autonomous Systems | | |
|---|---|---|
| Step 1 | Configure ECMP. | See Configure ECMP on a Virtual Router. |
| Step 2 | For BGP routing, enable ECMP over multiple autonomous systems. | 1. Select **Network > Virtual Routers** and select the virtual router on which to enable ECMP for multiple BGP autonomous systems.<br>2. Select **BGP > Advanced** and select the **ECMP Multiple AS Support** check box. |
| Step 3 | Save the configuration. | Click **OK** and **Commit** the configuration. |

 © Palo Alto Networks, Inc.

## Verify ECMP

A virtual router configured for ECMP indicates in the Forwarding Information Base (FIB) table which routes are ECMP routes. An ECMP flag (E) for a route indicates that it is participating in ECMP for the egress interface to the next hop for that route.

| **Confirm That Routes Are Equal-Cost Multiple Paths** | |
|---|---|
| Look at the FIB and confirm that some routes are equal-cost multiple paths. | 1. Select **Network > Virtual Routers**.<br>2. In the row of the virtual router for which you enabled ECMP, click **More Runtime Stats**.<br>3. Select **Routing** > **Forwarding Table** to see the FIB. In the table, note that multiple routes to the same Destination (out a different Interface) have the "E" flag.<br>An asterisk "*" denotes the preferred path for the ECMP group. |

# LLDP

Palo Alto Networks firewalls support Link Layer Discovery Protocol (LLDP), which functions at the link layer to discover neighboring devices and their capabilities. LLDP allows the firewall and other network devices to send and receive LLDP data units (LLDPDUs) to and from neighbors. The receiving device stores the information in a MIB, which the Simple Network Management Protocol (SNMP) can access. LLDP makes troubleshooting easier, especially for virtual wire deployments where the firewall would typically go undetected by a ping or traceroute.


- LLDP Overview
- Supported TLVs in LLDP
- LLDP Syslog Messages and SNMP Traps
- Configure LLDP
- View LLDP Settings and Status
- Clear LLDP Statistics


 © Palo Alto Networks, Inc.

## LLDP Overview

LLDP operates at Layer 2 of the OSI model, using MAC addresses. An LLDPDU is a sequence of type-length-value (TLV) elements encapsulated in an Ethernet frame. The IEEE 802.1AB standard defines three MAC addresses for LLDPDUs: 01-80-C2-00-00-0E, 01-80-C2-00-00-03, and 01-80-C2-00-00-00.

The Palo Alto Networks firewall supports only one MAC address for transmitting and receiving LLDP data units: 01-80-C2-00-00-0E. When transmitting, the firewall uses 01-80-C2-00-00-0E as the destination MAC address. When receiving, the firewall processes datagrams with 01-80-C2-00-00-0E as the destination MAC address. If the firewall receives either of the other two MAC addresses for LLDPDUs on its interfaces, the firewall takes the same forwarding action it took prior to this feature, as follows:

- If the interface type is vwire, the firewall forwards the datagram to the other port.
- If the interface type is L2, the firewall floods the datagram to the rest of the VLAN.
- If the interface type is L3, the firewall drops the datagrams.

The PA-2000 Series platform is not supported due to the hardware limitation of how Aggregated Ethernet interfaces function. Panorama, the GlobalProtect Mobile Security Manager, and the WildFire appliance are also not supported.

Interface types that do not support LLDP are TAP, high availability (HA), Decrypt Mirror, virtual wire/vlan/L3 subinterfaces, and PA-7000 Series Log Processing Card (LPC) interfaces.

An LLDP Ethernet frame has the following format:

| Preamble | Destination MAC | Source MAC | Ethertype | Chassis ID TLV | Port ID TLV | Time To Live TLV | Optional TLVs | End of LLDPDU TLV | Frame Check Sequence |
|---|---|---|---|---|---|---|---|---|---|
| | 01:80:C2:00:00:0E or 01:80:C2:00:00:03 or 01:80:C2:00:00:00 | Station's Address | 0x88CC | Type=1 | Type=2 | Type=3 | Zero or more complete TLVs | Type=0, Length=0 | |

Within the LLDP Ethernet frame, the TLV structure has the following format:

| TLV Type | TLV Information String Length | TLV Information String |
|---|---|---|
| 7 bits | 9 bits | 0-511 octets |

## Supported TLVs in LLDP

LLDPDUs include mandatory and optional TLVs. The following table lists the mandatory TLVs that the firewall supports:

**Mandatory TLVs in an LLDPDU Message**

| Mandatory TLVs | TLV Type | Description |
|---|---|---|
| Chassis ID TLV | 1 | Identifies the firewall chassis. Each firewall must have exactly one unique Chassis ID. The Chassis ID subtype is 4 (MAC address) on Palo Alto Networks platforms will use the MAC address of Eth0 to ensure uniqueness. |
| Port ID TLV | 2 | Identifies the port from which the LLDPDU is sent. Each firewall uses one Port ID for each LLDPDU message transmitted. The Port ID subtype is 5 (interface name) and uniquely identifies the transmitting port. The firewall uses the interface's ifname as the Port ID. |
| Time-to-live (TTL) TLV | 3 | Specifies how long (in seconds) LLDPDU information received from the peer is retained as valid in the local firewall (range is 0-65535). The value is a multiple of the LLDP Hold Time Multiplier. When the TTL value is 0, the information associated with the device is no longer valid and the firewall removes that entry from the MIB. |
| End of LLDPDU TLV | 0 | Indicates the end of the TLVs in the LLDP Ethernet frame. |

The following table lists the optional TLVs that the Palo Alto Networks firewall supports:

**Optional TLVs in an LLDPDU Message**

| Optional TLVs | TLV Type | Purpose and Notes Regarding Firewall Implementation |
|---|---|---|
| Port Description TLV | 4 | Describes the port of the firewall in alpha-numeric format. The ifAlias object is used. |
| System Name TLV | 5 | Configured name of the firewall in alpha-numeric format. The sysName object is used. |
| System Description TLV | 6 | Describes the firewall in alpha-numeric format. The sysDescr object is used. |
| System Capabilities | 7 | Describes the deployment mode of the interface, as follows:<br>• An L3 interface is advertised with router (bit 6) capability and the "other" bit (bit 1).<br>• An L2 interface is advertised with MAC Bridge (bit 3) capability and the "other" bit (bit 1).<br>• A virtual wire interface is advertised with Repeater (bit 2) capability and the "other" bit (bit 1). |


© Palo Alto Networks, Inc.

**Optional TLVs in an LLDPDU Message (Continued)**

| Optional TLVs | TLV Type | Purpose and Notes Regarding Firewall Implementation |
|---|---|---|
| Management Address | 8 | One or more IP addresses used for the management of the device, as follows:<br>• IP address of the management (MGT) interface<br>• IPv4 and/or IPv6 address of the interface<br>• Loopback address<br>• User-defined address entered in the management address field<br>If no management IP address is provided, the default is the MAC address of the transmitting interface.<br>Included is the interface number of the management address specified. Also included is the OID of the hardware interface with the management address specified (if applicable).<br>If more than one management address is specified, they will be sent in the order they are specified, starting at the top of the list. A maximum of four Management Addresses are supported.<br>This is an optional parameter and can be left disabled. |

## LLDP Syslog Messages and SNMP Traps

The firewall stores LLDP information in MIBs, which an SNMP Manager can monitor. If you want the firewall to send SNMP trap notifications and syslog messages about LLDP events, you must enable **SNMP Syslog Notification** in an LLDP profile.

Per RFC 5424, The Syslog Protocol, and RFC 1157, A Simple Network Management Protocol, LLDP sends syslog and SNMP trap messages when MIB changes occur. These messages are rate-limited by the **Notification Interval**, an LLDP global setting that defaults to 5 seconds and is configurable.

Because the LLDP syslog and SNMP trap messages are rate-limited, some LLDP information provided to those processes might not match the current LLDP statistics seen when you View the LLDP status information. This is normal, expected behavior.

A maximum of 5 MIBs can be received per interface (Ethernet or AE). Each different source has one MIB. If this limit is exceeded, the error message `tooManyNeighbors` is triggered.

 © Palo Alto Networks, Inc.

# Configure LLDP

To configure LLDP, and create an LLDP profile, you must be a superuser or device administrator (deviceadmin). A firewall interface supports a maximum of five LLDP peers.

| Configure LLDP | | |
|---|---|---|
| Step 1 | Enable LLDP on the firewall. | Select **Network > LLDP** and edit the LLDP General section; select the **Enable** check box. |
| Step 2 | (Optional) Change LLDP global settings. | 1. For **Transmit Interval (sec)**, specify the interval (in seconds) at which LLDPDUs are transmitted. Default: 30 seconds. Range: 1-3600 seconds.<br>2. For **Transmit Delay (sec)**, specify the delay time (in seconds) between LLDP transmissions sent after a change is made in a TLV element. The delay helps to prevent flooding the segment with LLDPDUs if many network changes spike the number of LLDP changes, or if the interface flaps. The **Transmit Delay** must be less than the **Transmit Interval**. Default: 2 seconds. Range: 1-600 seconds.<br>3. For **Hold Time Multiple**, specify a value that is multiplied by the **Transmit Interval** to determine the total TTL Hold Time. Default: 4. Range: 1-100. The maximum TTL Hold Time is 65535 seconds, regardless of the multiplier value.<br>4. For **Notification Interval**, specify the interval (in seconds) at which LLDP Syslog Messages and SNMP Traps are transmitted when MIB changes occur. Default: 5 seconds. Range: 1-3600 seconds.<br>5. Click **OK**. |

| Configure LLDP | |
|---|---|
| Step 3 Create an LLDP profile.<br>For descriptions of the optional TLVs, see Supported TLVs in LLDP. | 1. Select **Network > Network Profiles > LLDP Profile** and click **Add**.<br>2. Enter a **Name** for the LLDP profile.<br>3. For **Mode**, select **transmit-receive** (default), **transmit-only**, or **receive-only**.<br>4. Click the **SNMP Syslog Notification** check box to enable SNMP notifications and syslog messages. If enabled, the global **Notification Interval** is used. The firewall will send both an SNMP trap and a syslog event as configured in the **Device > Log Settings > System > SNMP Trap Profile** and **Syslog Profile**.<br>5. For Optional TLVs, select the TLVs you want transmitted:<br>• **Port Description**<br>• **System Name**<br>• **System Description**<br>• **System Capabilities**<br>6. Specifying a **Management Address** is optional. To add one or more, select the check box and **Add** a **Name**.<br>7. Select the **Interface** from which to obtain the management address. At least one management address is required if **Management Address** TLV is enabled. If no management IP address is configured, the system uses the MAC address of the transmitting interface as the management address TLV.<br>8. Select **IPv4** or **IPv6**, and in the adjacent field, select an IP address from the drop-down (which lists the addresses configured on the selected interface), or enter an address.<br>9. Click **OK**.<br>10. Up to four management addresses are allowed. If you specify more than one **Management Address**, they will be sent in the order they are specified, starting at the top of the list. To change the order of the addresses, select an address and use the **Move Up** or **Move Down** buttons.<br>11. Click **OK**. |
| Step 4 Assign an LLDP profile to an interface. | 1. Select **Network > Interfaces** and select the interface where you will assign an LLDP profile.<br>2. Select **Advanced > LLDP**.<br>3. Select the **Enable LLDP** check box to assign an LLDP profile to the interface.<br>4. For **Profile**, select the profile you created. Selecting **None** enables LLDP with basic functionality: sends the three mandatory TLVs and enables **transmit-receive** mode.<br>If you want to create a new profile, click **LLDP Profile** and follow the instructions in Step 4.<br>5. Click **OK**. |
| Step 5 Save the configuration. | Click **Commit**. |


© Palo Alto Networks, Inc.

## View LLDP Settings and Status

Perform the following procedure to view LLDP settings and status.

| View LLDP Settings and Status | |
|---|---|
| Step 1 View LLDP global settings. | 1. Select **Network > LLDP**.<br>• On the LLDP General screen, the **Enable** check box indicates whether LLDP is enabled or not.<br>– If LLDP is enabled, the configured global settings (Transmit Interval, Transmit Delay, Hold Time Multiple, and Notification Interval) are displayed.<br>– If LLDP is not enabled, the default values of the global settings are displayed.<br>For descriptions of these values, see (Optional) Change LLDP global settings. |

| View LLDP Settings and Status (Continued) | |
|---|---|
| Step 2 View the LLDP status information. | 1. Select the **Status** tab.<br>2. (Optional) Enter a filter to restrict the information that is displayed.<br>Interface Information:<br>• **Interface**—Name of the interfaces that have LLDP profiles assigned to them.<br>• **LLDP**—LLDP status: enabled or disabled.<br>• **Mode**— LLDP mode of the interface: Tx/Rx, Tx Only, or Rx Only.<br>• **Profile**—Name of the profile assigned to the interface.<br>Transmission Information:<br>• **Total Transmitted**—Count of LLDPDUs transmitted out the interface.<br>• **Dropped Transmit**—Count of LLDPDUs that were not transmitted out the interface because of an error. For example, a length error when the system is constructing an LLDPDU for transmission.<br>Received Information:<br>• **Total Received**—Count of LLDP frames received on the interface.<br>• **Dropped TLV**—Count of LLDP frames discarded upon receipt.<br>• **Errors**—Count of TLVs that were received on the interface and contained errors. Types of TLV errors include: one or more mandatory TLVs missing, out of order, containing out-of-range information, or length error.<br>• **Unrecognized**—Count of TLVs received on the interface that are not recognized by the LLDP local agent. For example, the TLV type is in the reserved TLV range.<br>• **Aged Out**—Count of items deleted from the Receive MIB due to proper TTL expiration. |


© Palo Alto Networks, Inc.

| View LLDP Settings and Status (Continued) | |
|---|---|
| Step 3 View summary LLDP information for each neighbor seen on an interface. | 1. Select the **Peers** tab.<br>2. Optionally enter a filter to restrict the information being displayed.<br>• Local Interface—Interface on the firewall that detected the neighboring device.<br>• Remote Chassis ID—Chassis ID of the peer. The MAC address will be used.<br>• Port ID—Port ID of the peer.<br>• Name—Name of peer.<br>• More info—Provides the following remote peer details, which are based on the Mandatory and Optional TLVs:<br>– Chassis Type: MAC address.<br>– MAC Address: MAC address of the peer.<br>– System Name: Name of the peer.<br>– System Description: Description of the peer.<br>– Port Description: Port description of the peer.<br>– Port Type: Interface name.<br>– Port ID: The firewall uses the interface's ifname.<br>– System Capabilities: Capabilities of the system. O=Other, P=Repeater, B=Bridge, W=Wireless-LAN, R=Router, T=Telephone<br>– Enabled Capabilities: Capabilities enabled on the peer.<br>– Management Address: Management address of the peer. |

## Clear LLDP Statistics

You can clear LLDP statistics for specific interfaces.

| Clear LLDP Statistics | | |
|---|---|---|
| Step 1 | Clear LLDP statistics for specific interfaces. | 1. Select **Network > LLDP > Status** and in the left hand column, select one or more interfaces for which you want to clear LLDP statistics.<br>2. Click **Clear LLDP Statistics** at the bottom of the screen. |


© Palo Alto Networks, Inc.

# Policy

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Captive Portal, Denial of Service (DoS), and Zone protection policies. All these different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. The following topics describe how to work with policy:

- Policy Types
- Security Policy
- Policy Objects
- Security Profiles
- Enumeration of Rules Within a Rulebase
- Move or Clone a Policy Rule or Object to a Different Virtual System
- Use Tags to Group and Visually Distinguish Objects
- Use a Dynamic Block List in Policy
- Register IP Addresses and Tags Dynamically
- Monitor Changes in the Virtual Environment
- CLI Commands for Dynamic IP Addresses and Tags
- Identify Users Connected through a Proxy Server
- Policy-Based Forwarding
- DoS Protection Against Flooding of New Sessions

© Palo Alto Networks, Inc.

# Policy Types

The Palo Alto Networks next-generation firewall supports a variety of policy types that work together to safely enable applications on your network.

| Policy Type | Description |
|---|---|
| Security | Determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service. For more details, see Security Policy. |
| NAT | Instruct the firewall which packets need translation and how to do the translation. The firewall supports both source address and/or port translation and destination address and/or port translation. For more details, see NAT. |
| QoS | Identify traffic requiring QoS treatment (either preferential treatment or bandwidth-limiting) using a defined parameter or multiple parameters and assign it a class. For more details, see Quality of Service. |
| Policy Based Forwarding | Identify traffic that should use a different egress interface than the one that would normally be used based on the routing table. For details, see Policy-Based Forwarding. |
| Decryption | Identify encrypted traffic that you want to inspect for visibility, control, and granular security. For more details, see Decryption. |
| Application Override | Identify sessions that you do not want processed by the App-ID engine, which is a Layer-7 inspection. Traffic matching an application override policy forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4. For more details, see Manage Custom or Unknown Applications. |
| Captive Portal | Identify traffic that requires the user to be known. The captive portal policy is only triggered if other User-ID mechanisms did not identify a user to associate with the source IP address. For more details, see Captive Portal. |
| DoS Protection | Identify potential denial-of-service (DoS) attacks and take protective action in response to rule matches. DoS Protection Profiles. |


© Palo Alto Networks, Inc.

# Security Policy

Security policies protect network assets from threats and disruptions and aid in optimally allocating network resources for enhancing productivity and efficiency in business processes. On the Palo Alto Networks firewall, security policies determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service.

All traffic passing through the firewall is matched against a session and each session is matched against a security policy. When a session match occurs, the security policy is applied to bi-directional traffic (client to server and server to client) in that session. For traffic that doesn't match any defined rules, the default rules apply. The default rules—displayed at the bottom of the security rulebase—are predefined to allow all intrazone (within the zone) traffic and deny all interzone (between zones) traffic. Although these rules are part of the pre-defined configuration and are read-only by default, you can override them and change a limited number of settings, including the tags, action (allow or block), log settings, and security profiles.

Security policies are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log, if logging is enabled for that rule. The logging options are configurable for each rule, and can for example be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.


- Components of a Security Policy Rule
- Security Policy Best Practices

# Components of a Security Policy Rule

The security policy rule construct permits a combination of the required and optional fields as detailed in the following tables:

- Required Fields
- Optional Fields

## Required Fields

| Required Field | Description |
|---|---|
| **Name** | A label that supports up to 31 characters, used to identify the rule. |
| **Rule Type** | Specifies whether the rule applies to traffic within a zone, between zones, or both:<br>• **universal** (default)—Applies the rule to all matching interzone and intrazone traffic in the specified source and destination zones. For example, if you create a universal role with source zones A and B and destination zones A and B, the rule would apply to all traffic within zone A, all traffic within zone B, and all traffic from zone A to zone B and all traffic from zone B to zone A.<br>• **intrazone**—Applies the rule to all matching traffic within the specified source zones (you cannot specify a destination zone for intrazone rules). For example, if you set the source zone to A and B, the rule would apply to all traffic within zone A and all traffic within zone B, but not to traffic between zones A and B.<br>• **interzone**—Applies the rule to all matching traffic between the specified source and destination zones. For example, if you set the source zone to A, B, and C and the destination zone to A and B, the rule would apply to traffic from zone A to zone B, from zone B to zone A, from zone C to zone A, and from zone C to zone B, but not traffic within zones A, B, or C. |
| **Source Zone** | The zone from which the traffic originates. |
| **Destination Zone** | The zone at which the traffic terminates. If you use NAT, make sure to always reference the post-NAT zone. |
| **Application** | The application which you wish to control. The firewall uses App-ID, the traffic classification technology, to identify traffic on your network. App-ID provides application control and visibility in creating security policies that block unknown applications, while enabling, inspecting, and shaping those that are allowed. |

 © Palo Alto Networks, Inc.

| Required Field | Description (Continued) |
|---|---|
| **Action** | Specifies an *Allow* or *Block* action for the traffic based on the criteria you define in the rule. When you configure the firewall to block traffic, it either resets the connection or silently drops packets. To provide a better user experience, you can configure granular options to block traffic instead of silently dropping packets, which can cause some applications to break and appear unresponsive to the user.<br>• **Allow**—(default action) Allows the traffic.<br>• **Deny**—Blocks traffic, and enforces the default *Deny Action* defined for the application that is being denied. To view the deny action defined by default for an application, view the application details in **Objects > Applications** or check the application details in Applipedia.<br>• **Drop**—Silently drops the traffic; for an application, it overrides the default deny action. A TCP reset is not sent to the host/application.<br>For Layer 3 interfaces, to optionally send an ICMP unreachable response to the client, set Action: **Drop** and enable the **Send ICMP Unreachable** checkbox. When enabled, the firewall sends the ICMP code for *communication with the destination is administratively prohibited*— ICMPv4: Type 3, Code 13; ICMPv6: Type 1, Code 1.<br>• Reset—Resets a connection in one of the following ways.<br>  • **Reset client**—Sends a TCP reset to the client-side device.<br>  • **Reset server**—Sends a TCP reset to the server-side device.<br>  • **Reset both**—Sends a TCP reset to both the client-side and server-side devices.<br>A reset is sent only after a session is formed. If the session is blocked before a 3-way handshake is completed, the firewall does not send a reset. For a TCP session with a reset action, the firewall does not send an ICMP Unreachable response. For a UDP session with a drop or reset action, if the **ICMP Unreachable** checkbox is selected, the firewall sends an ICMP message to the client. |

## Optional Fields

| Optional Field | Description |
|---|---|
| **Tag** | A keyword or phrase that allows you to filter security rules. This is handy when you have defined many rules and wish to then review those that are tagged with a keyword such as *IT-sanctioned applications* or *High-risk applications*. |
| **Description** | A text field, up to 255 characters, used to describe the rule. |
| **Source IP Address** | Define host IP or FQDN, subnet, named groups, or country-based enforcement. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address). |
| **Destination IP Address** | The location or destination for the traffic. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address). |
| **User** | The user or group of users for whom the policy applies. You must have User-ID enabled on the zone. To enable User-ID, see User-ID Overview. |

| Optional Field | Description (Continued) |
|---|---|
| **URL Category** | Using the URL Category as match criteria allows you to customize security profiles (antivirus, anti-spyware, vulnerability, file-blocking, Data Filtering, and DoS) on a per-URL-category basis. For example, you can prevent.exe file download/upload for URL categories that represent higher risk while allowing them for other categories. This functionality also allows you to attach schedules to specific URL categories (allow social-media websites during lunch & after-hours), mark certain URL categories with QoS (financial, medical, and business), and select different log forwarding profiles on a per-URL-category-basis.<br>Although you can manually configure URL categories on your device, to take advantage of the dynamic URL categorization updates available on the Palo Alto Networks firewalls, you must purchase a URL filtering license.<br> To block or allow traffic based on URL category, you must apply a URL Filtering profile to the security policy rules. Define the URL Category as *Any* and attach a URL Filtering profile to the security policy. See Define Basic Security Rules for information on using the default profiles in your security policy and see Control Access to Web Content for more details. |
| **Service** | Allows you to select a Layer 4 (TCP or UDP) port for the application. You can choose *any*, specify a port, or use *application-default* to permit use of the standards-based port for the application. For example, for applications with well- known port numbers such as DNS, the *application-default* option will match against DNS traffic only on TCP port 53. You can also add a custom application and define the ports that the application can use.<br> For inbound allow rules (for example, from untrust to trust), using *application-default* prevents applications from running on unusual ports and protocols. Application-default is the default option; while the device still checks for all applications on all ports, with this configuration, applications are only allowed on their standard ports/protocols. |
| **Security Profiles** | Provide additional protection from threats, vulnerabilities, and data leaks. Security profiles are only evaluated for rules that have an *allow* action. |
| **HIP Profile** (for GlobalProtect) | Allows you to identify clients with Host Information Profile (HIP) and then enforce access privileges. |
| **Options** | Allow you to define logging for the session, log forwarding settings, change Quality of Service (QoS) markings for packets that match the rule, and schedule when (day and time) the security rule should be in effect. |

 © Palo Alto Networks, Inc.

## Security Policy Best Practices

The task of safely enabling Internet access and preventing misuse of web access privileges, and exposure to vulnerabilities and attacks is a continuous process. The key principle when defining policy on the Palo Alto Networks firewall is to use a positive enforcement approach. Positive enforcement implies that you selectively allow what is required for day-to-day business operations as opposed to a negative enforcement approach where you would selectively block everything that is not allowed. Consider the following suggestions when creating policy:

- If you have two or more zones with identical security requirements, combine them into one security rule.
- The ordering of rules is crucial to ensure the best match criteria. Because policy is evaluated top down, the more specific policy must precede the ones that are more general, so that the more specific rule is not *shadowed*. The term shadow refers to a rule that is not evaluated or is skipped because it is placed lower in the policy list. When the rule is placed lower, it is not evaluated because the match criteria was met by another rule that preceded it, thereby shadowing the rule from policy evaluation.
- To restrict and control access to inbound applications, in the security policy, explicitly define the port that the service/application will be listening on.
- Logging for broad allow rules—for example access to well known servers like DNS—can generate a lot of traffic. Hence it is not recommended unless absolutely necessary.
- By default, the firewall creates a log entry at the end of a session. However, you can modify this default behavior and configure the firewall to log at the start of the session. Because this significantly increases the log volume, logging at session start is recommended only when you are troubleshooting an issue. Another alternative for troubleshooting without enabling logging at session start is to use the session browser (**Monitor > Session Browser**) to view the sessions in real time.

© Palo Alto Networks, Inc.

# Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With policy objects that are a collective unit, you can reference the object in security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

You can create the following policy objects on the firewall:

| Policy Object | Description |
|---|---|
| Address/Address Group, Region | Allow you to group specific source or destination addresses that require the same policy enforcement. The address object can include an IPv4 or IPv6 address (single IP, range, subnet) or the FQDN. Alternatively, a region can be defined by the latitude and longitude coordinates or you can select a country and define an IP address or IP range. You can then group a collection of address objects to create an *address group* object.<br>You can also use dynamic address groups to dynamically update IP addresses in environments where host IP addresses change frequently. |
| User/User Group | Allow you to create a list of users from the local database or an external database and group them. |
| Application Group and Application Filter | An *Application Filter* allows you to filter applications dynamically. It allows you to filter, and save a group of applications using the attributes defined in the application database on the firewall. For example, you can Create an Application Filter by one or more attributes—category, sub-category, technology, risk, characteristics. With an application filter, when a content update occurs, any new applications that match your filter criteria are automatically added to your saved application filter.<br>An *Application Group* allows you to create a static group of specific applications that you want to group together for a group of users or for a particular service, or to achieve a particular policy goal. See Create an Application Group. |
| Service/Service Groups | Allows you to specify the source and destination ports and protocol that a service can use. The firewall includes two pre-defined services—service-http and service-https—that use TCP ports 80 and 8080 for HTTP, and TCP port 443 for HTTPS. You can however, create any custom service on any TCP/UDP port of your choice to restrict application usage to specific ports on your network (in other words, you can define the default port for the application).<br> To view the standard ports used by an application, in **Objects > Applications** search for the application and click the link. A succinct description displays. |

 © Palo Alto Networks, Inc.

# Security Profiles

While security policies enable you to allow or block traffic on your network, security profiles help you define an *allow but scan* rule, which scan allowed applications for threats, such as viruses, malware, spyware, and DDOS attacks. When traffic matches the allow rule defined in the security policy, the security profile(s) that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.

Security profiles are not used in the match criteria of a traffic flow. The security profile is applied to scan traffic after the application or category is allowed by the security policy.

The firewall provides default security profiles that you can use out of the box to begin protecting your network from threats. See Set Up Basic Security Policies for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, you can create custom profiles. See Scan Traffic for Threats for more information.

You can add security profiles that are commonly applied together to a security profile group; this set of profiles can be treated as a unit and added to security policies in one step (or included in security policies by default, if you choose to set up a default security profile group).

The following topics provide more detailed information about each type of security profile and how to set up a security profile group:

- Antivirus Profiles
- Anti-Spyware Profiles
- Vulnerability Protection Profiles
- URL Filtering Profiles
- Data Filtering Profiles
- File Blocking Profiles
- WildFire Analysis Profiles
- DoS Protection Profiles
- Zone Protection Profiles
- Security Profile Group

© Palo Alto Networks, Inc.

## Antivirus Profiles

Antivirus profiles protect against viruses, worms, and trojans as well as spyware downloads. Using a stream-based malware prevention engine, which inspects traffic the moment the first packet is received, the Palo Alto Networks antivirus solution can provide protection for clients without significantly impacting the performance of the firewall. This profile scans for a wide variety of malware in executables, PDF files, HTML and JavaScript viruses, including support for scanning inside compressed files and data encoding schemes. If you have enabled Decryption on the firewall, the profile also enables scanning of decrypted content.

The default profile inspects all of the listed protocol decoders for viruses, and generates alerts for SMTP, IMAP, and POP3 protocols while blocking for FTP, HTTP, and SMB protocols. You can configure the action for a decoder or antivirus signature and specify how the firewall responds to a threat event:

- **Default**—For each threat signature and Antivirus signature that is defined by Palo Alto Networks, a default action is specified internally. Typically, the default action is an alert or a reset-both. The default action is displayed in parenthesis, for example default (alert) in the threat or Antivirus signature.
- **Allow**—Permits the application traffic
- **Alert**—Generates an alert for each application traffic flow. The alert is saved in the threat log.
- **Drop**—Drops the application traffic.
- **Reset Client**—For TCP, resets the client-side connection. For UDP, drops the connection.
- **Reset Server**—For TCP, resets the server-side connection. For UDP, drops the connection.
- **Reset Both**—For TCP, resets the connection on both client and server ends. For UDP, drops the connection.

Customized profiles can be used to minimize antivirus inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms.

The Palo Alto Networks WildFire system also provides signatures for persistent threats that are more evasive and have not yet been discovered by other antivirus solutions. As threats are discovered by WildFire, signatures are quickly created and then integrated into the standard antivirus signatures that can be downloaded by Threat Prevention subscribers on a daily basis (sub-hourly for WildFire subscribers).

 © Palo Alto Networks, Inc.

## Anti-Spyware Profiles

Anti-Spyware profiles blocks spyware on compromised hosts from trying to phone-home or beacon out to external command-and-control (C2) servers, allowing you to detect malicious traffic leaving the network from infected clients. You can apply various levels of protection between zones. For example, you may want to have custom Anti-Spyware profiles that minimize inspection between trusted zones, while maximizing inspection on traffic received from an untrusted zone, such as Internet facing zones.

You can define your own custom Anti-Spyware profiles, or choose one of the following predefined profiles when applying anti-spyware to a security policy:

- **Default**—Uses the default action for every signature, as specified by Palo Alto Networks when the signature is created.
- **Strict**—Overrides the default action of critical, high, and medium severity threats to the block action, regardless of the action defined in the signature file. This profile still uses the default action for medium and informational severity signatures.

When the firewall detects a threat event, you can configure the following actions in an Anti-Spyware profile:

- **Default**—For each threat signature and anti-spyware signature that is defined by Palo Alto Networks, a default action is specified internally. Typically the default action is an alert or a reset-both. The default action is displayed in parenthesis, for example default (alert) in the threat or Antivirus signature.
- **Allow**—Permits the application traffic
- **Alert**—Generates an alert for each application traffic flow. The alert is saved in the threat log.
- **Drop**—Drops the application traffic.
- **Reset Client**—For TCP, resets the client-side connection. For UDP, drops the connection.
- **Reset Server**—For TCP, resets the server-side connection. For UDP, drops the connection.
- **Reset Both**—For TCP, resets the connection on both client and server ends. For UDP, drops the connection.
- **Block IP**— This action blocks traffic from either a source or a source-destination pair. It is configurable for a specified period of time.

In addition, you can enable the DNS Sinkholing action in Anti-Spyware profiles to enable the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. This feature helps to identify infected hosts on the protected network using DNS traffic Infected hosts can then be easily identified in the traffic and threat logs because any host that attempts to connect to the sinkhole IP address are most likely infected with malware.

Anti-Spyware and Vulnerability Protection profiles are configured similarly.

© Palo Alto Networks, Inc.

## Vulnerability Protection Profiles

Vulnerability Protection profiles stop attempts to exploit system flaws or gain unauthorized access to systems. While Anti-Spyware profiles help identify infected hosts as traffic leaves the network, Vulnerability Protection profiles protect against threats entering the network. For example, Vulnerability Protection profiles help protect against buffer overflows, illegal code execution, and other attempts to exploit system vulnerabilities. The default Vulnerability Protection profile protects clients and servers from all known critical, high, and medium-severity threats. You can also create exceptions, which allow you to change the response to a specific signature.

To configure how the firewall responds to a threat, see Anti-Spyware Profiles for a list of supported actions.

 © Palo Alto Networks, Inc.

## URL Filtering Profiles

URL Filtering profiles enable you to monitor and control how users access the web over HTTP and HTTPS. The firewall comes with a default profile that is configured to block websites such as known malware sites, phishing sites, and adult content sites. You can use the default profile in a security policy, clone it to be used as a starting point for new URL filtering profiles, or add a new URL profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories.

# Data Filtering Profiles

Data filtering profiles prevent sensitive information such as credit card or social security numbers from leaving a protected network. The data filtering profile also allows you to filter on key words, such as a sensitive project name or the word confidential. It is important to focus your profile on the desired file types to reduce false positives. For example, you may only want to search Word documents or Excel spreadsheets. You may also only want to scan web-browsing traffic, or FTP.

You can use default profiles, or create custom data patterns. There are two default profiles:

- CC# (Credit Card)—Identifies credit card numbers using a hash algorithm. The content must match the hash algorithm in order for data to be detected as a credit card number. This method will reduce false positives.
- SSN# (Social Security Number)—Uses an algorithm to detect nine digit numbers, regardless of format. There are two fields: SSN# and SSN# (no dash).

## Weight and Threshold Values

It is important to understand how the weight of an object (SSN, CC#, pattern) is calculated in order to set the appropriate threshold for a condition you are trying to filter. Each occurrence multiplied by the weight value will be added together in order to reach an action threshold (alert or block).

### Example: Filter for Social Security Numbers Only

For simplicity, if you only want to filter files with Social Security Numbers (SSN) and you define a weight of 3 for SSN#, you would use the following formula: each instance of a SSN x weight = threshold increment. In this case, if a Word document has 10 social security numbers you multiply that by the weight of 3, so 10 x 3 = 30. In order to take action for a file that contains 10 social security numbers you would set the threshold to 30. You may want to set an alert at 30 and then block at 60. You may also want to set a weight in the field SSN# (no dash) for Social Security Numbers that do not contain dashes. If multiple settings are used, they will accumulate to reach a given threshold.

### Example: Filter for Social Security Numbers and a Custom Pattern

In this example, we will filter on files that contain Social Security Numbers and the custom pattern confidential. In other words, if a file has Social Security Numbers in addition to the word confidential and the combined instances of those items hit the threshold, the file will trigger an alert or block, depending on the action setting.

SSN# weight = 3

Custom Pattern confidential weight = 20

The custom pattern is case sensitive.

 © Palo Alto Networks, Inc.

If the file contains 20 Social Security Numbers and a weight of 3 is configured, that is 20 x 3 = 60. If the file also contains one instance of the term confidential and a weight of 20 is configured, that is 1 x 20 = 20 for a total of 80. If your threshold for block is set to 80, this scenario would block the file. The alert or block action will be triggered as soon as the threshold is hit.

## File Blocking Profiles

The firewall uses file blocking profiles to block specified file types over specified applications and in the specified session flow direction (inbound/outbound/both). You can set the profile to alert or block on upload and/or download and you can specify which applications will be subject to the file blocking profile. You can also configure custom block pages that will appear when a user attempts to download the specified file type. This allows the user to take a moment to consider whether or not they want to download a file.

Configure a file blocking profile with the following actions:

- **Alert**—When the specified file type is detected, a log is generated in the data filtering log.
- **Block**—When the specified file type is detected, the file is blocked and a customizable block page is presented to the user. A log is also generated in the data filtering log.
- **Continue**—When the specified file type is detected, a customizable response page is presented to the user. The user can click through the page to download the file. A log is also generated in the data filtering log. Because this type of forwarding action requires user interaction, it is only applicable for web traffic.

 © Palo Alto Networks, Inc.

## WildFire Analysis Profiles

Use a WildFire analysis profile to enable the firewall to forward unknown files or email links for WildFire analysis. Specify files to be forwarded for analysis based on application, file type, and transmission direction (upload or download). Files or email links matched to the profile rule are forwarded either the WildFire public cloud or the WildFire private cloud (hosted with a WF-500 appliance), depending on the analysis location defined for the rule.

You can also use the WildFire analysis profiles to set up a Wildfire hybrid cloud deployment. If you are using a WildFire appliance to analyze sensitive files locally (such as PDFs), you can specify for less sensitive files types (such as PE files) or file types that are not supported for WildFire appliance analysis (such as APKs) to be analyzed by the WildFire public cloud. Using both the WildFire appliance and the WildFire cloud for analysis allows you to benefit from a prompt verdict for files that have already been processed by the cloud, and for files that are not supported for appliance analysis, and frees up the appliance capacity to process sensitive content.

© Palo Alto Networks, Inc.

## DoS Protection Profiles

DoS protection profiles provide detailed control for Denial of Service (DoS) protection policies. DoS policies allow you to control the number of sessions between interfaces, zones, addresses, and countries based on aggregate sessions or source and/or destination IP addresses. There are two DoS protection mechanisms that the Palo Alto Networks firewalls support.

- **Flood Protection**—Detects and prevents attacks where the network is flooded with packets resulting in too many half-open sessions and/or services being unable to respond to each request. In this case the source address of the attack is usually spoofed. See DoS Protection Against Flooding of New Sessions.
- **Resource Protection**— Detects and prevent session exhaustion attacks. In this type of attack, a large number of hosts (bots) are used to establish as many fully established sessions as possible to consume all of a system's resources.

You can enable both types of protection mechanisms in a single DoS protection profile.

The DoS profile is used to specify the type of action to take and details on matching criteria for the DoS policy. The DoS profile defines settings for SYN, UDP, and ICMP floods, can enable resource protect and defines the maximum number of concurrent connections. After you configure the DoS protection profile, you then attach it to a DoS policy.

When configuring DoS protection, it is important to analyze your environment in order to set the correct thresholds and due to some of the complexities of defining DoS protection policies, this guide will not go into detailed examples. For more information, refer to the Threat Prevention Tech Note.

 © Palo Alto Networks, Inc.

## Zone Protection Profiles

Zone protection profiles provide additional protection between specific network zones in order to protect the zones against attack. The profile must be applied to the entire zone, so it is important to carefully test the profiles in order to prevent issues that may arise with the normal traffic traversing the zones. When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session. For more information, refer to the Threat Prevention Tech Note.

© Palo Alto Networks, Inc.

# Security Profile Group

A security profile group is a set of security profiles that can be treated as a unit and then easily added to security policies. Profiles that are often assigned together can be added to profile groups to simplify the creation of security policies. You can also setup a default security profile group—new security policies will use the settings defined in the default profile group to check and control traffic that matches the security policy. Name a security profile group *default* to allow the profiles in that group to be added to new security policies by default. This allows you to consistently include your organization's preferred profile settings in new policies automatically, without having to manually add security profiles each time you create new rules.

The following sections show how to create a security profile group and how to enable a profile group to be used by default in new security policies:

- Create a Security Profile Group
- Set Up or Override a Default Security Profile Group

## Create a Security Profile Group

Use the following steps to create a security profile group and add it to a security policy.

| **Create a Security Profile Group** | |
|---|---|
| Step 1 Create a security profile group. | 1. Select **Objects > Security Profile Groups** and **Add** a new security profile group.<br>2. Give the profile group a descriptive **Name**, for example, Threats.<br>3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.<br>4. Add existing profiles to the group.<br> 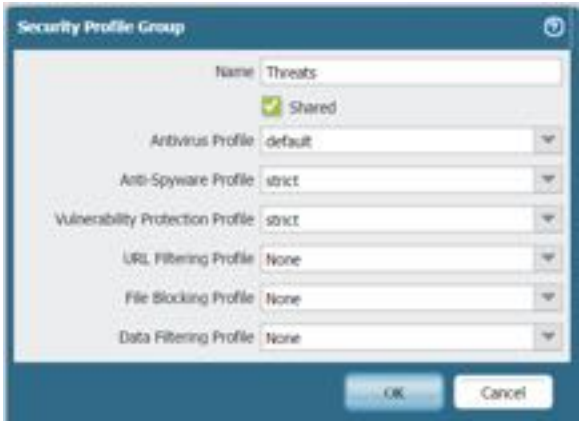 <br>5. Click **OK** to save the profile group. |

 © Palo Alto Networks, Inc.

| Create a Security Profile Group | | |
|---|---|---|
| Step 2 | Add a security profile group to a security policy. | 1. Select **Policies > Security** and **Add** or modify a security policy rule.<br>2. Select the **Actions** tab.<br>3. In the Profile Setting section, select **Group** for the **Profile Type**.<br>4. In the **Group Profile** drop-down, select the group you created (for example, select the Threats group):<br> 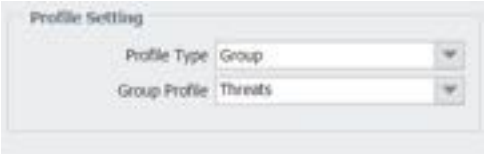 <br>5. Click **OK** to save the policy and **Commit** your changes. |
| Step 3 | Save your changes. | Click **Commit**. |

## Set Up or Override a Default Security Profile Group

Use the following options to set up a default security profile group to be used in new security policies, or to override an existing default group. When an administrator creates a new security policy, the default profile group will be automatically selected as the policy's profile settings, and traffic matching the policy will be checked according to the settings defined in the profile group (the administrator can choose to manually select different profile settings if desired). Use the following options to set up a default security profile group or to override your default settings.

If no default security profile exists, the profile settings for a new security policy are set to **None** by default.

| Set Up or Override a Default Security Profile Group | |
|---|---|
| • Create a security profile group. | 1. Select **Objects > Security Profile Groups** and Add a new security profile group.<br>2. Give the profile group a descriptive **Name**, for example, Threats.<br>3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.<br>4. Add existing profiles to the group. For details on creating profiles, see Security Profiles.<br> 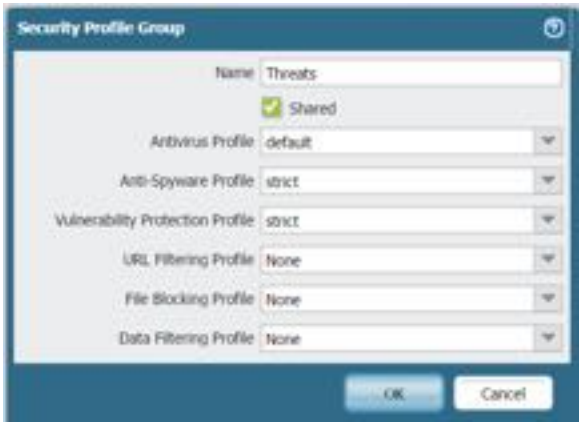 <br>5. Click **OK** to save the profile group.<br>6. **Add the security profile group to a security policy.**<br>7. **Add** or modify a security policy rule and select the **Actions** tab.<br>8. Select **Group** for the **Profile Type**.<br>9. In the **Group Profile** drop-down, select the group you created (for example, select the Threats group):<br> <br>10. Click **OK** to save the policy and **Commit** your changes. |

 © Palo Alto Networks, Inc.

| Set Up or Override a Default Security Profile Group | |
|---|---|
| • Set up a default security profile group. | 1. Select **Objects > Security Profile Groups** and add a new security profile group or modify an existing security profile group.<br>2. **Name** the security profile group *default*:<br> 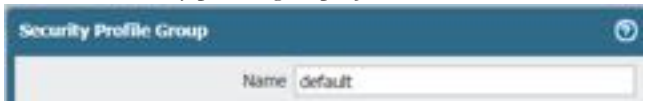 <br>3. Click **OK** and **Commit**.<br>4. Confirm that the *default* security profile group is included in new security policies by default:<br>a. Select **Policies > Security** and **Add** a new security policy.<br>b. Select the **Actions** tab and view the **Profile Setting** fields:<br> <br>By default, the new security policy correctly shows the **Profile Type** set to Group and the *default* **Group Profile** is selected. |
| • Override a default security profile group. | If you have an existing default security profile group, and you do not want that set of profiles to be attached to a new security policy, you can continue to modify the Profile Setting fields according to your preference. Begin by selecting a different Profile Type for your policy (**Policies > Security > Security Policy Rule > Actions**). |

# Enumeration of Rules Within a Rulebase

Each rule within a rulebase is automatically numbered and the ordering adjusts as rules are moved or reordered. When filtering rules to find rules that match the specified filter(s), each rule is listed with its number in the context of the complete set of rules in the rulebase and its place in the evaluation order.

On Panorama, pre-rules, post-rules, and default rules are independently numbered. When Panorama pushes rules to a firewall, the rule numbering reflects the hierarchy and evaluation order of shared rules, device-group pre-rules, firewall rules, device-group post-rules, and default rules. The **Preview Rules** option in Panorama displays an ordered list view of the total number of rules on a firewall.

### View the Ordered List of Rules Within a Rulebase

- View the numbered list of rules on the firewall.

  Select **Policies** and any rulebase under it. For example, **Policies > QoS**. The left-most column in the table displays the rule number.

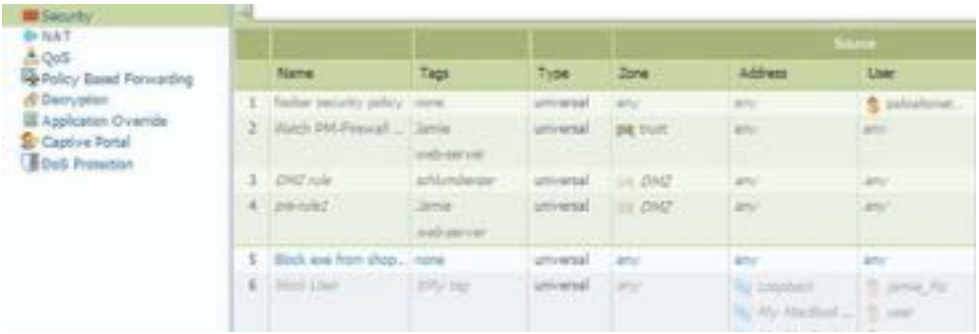


- View the numbered list of rules on Panorama.

  Select **Policies** and any rulebase under it. For example, **Policies > Security> Pre-rules**.

 © Palo Alto Networks, Inc.

**View the Ordered List of Rules Within a Rulebase (Continued)**

- After you push the rules from Panorama, view the complete list of rules with numbers on the firewall.

  From the web interface of the firewall, select **Policies** and pick any rulebase under it. For example, select **Policies > Security** and view the complete set of numbered rules that the firewall will evaluate.

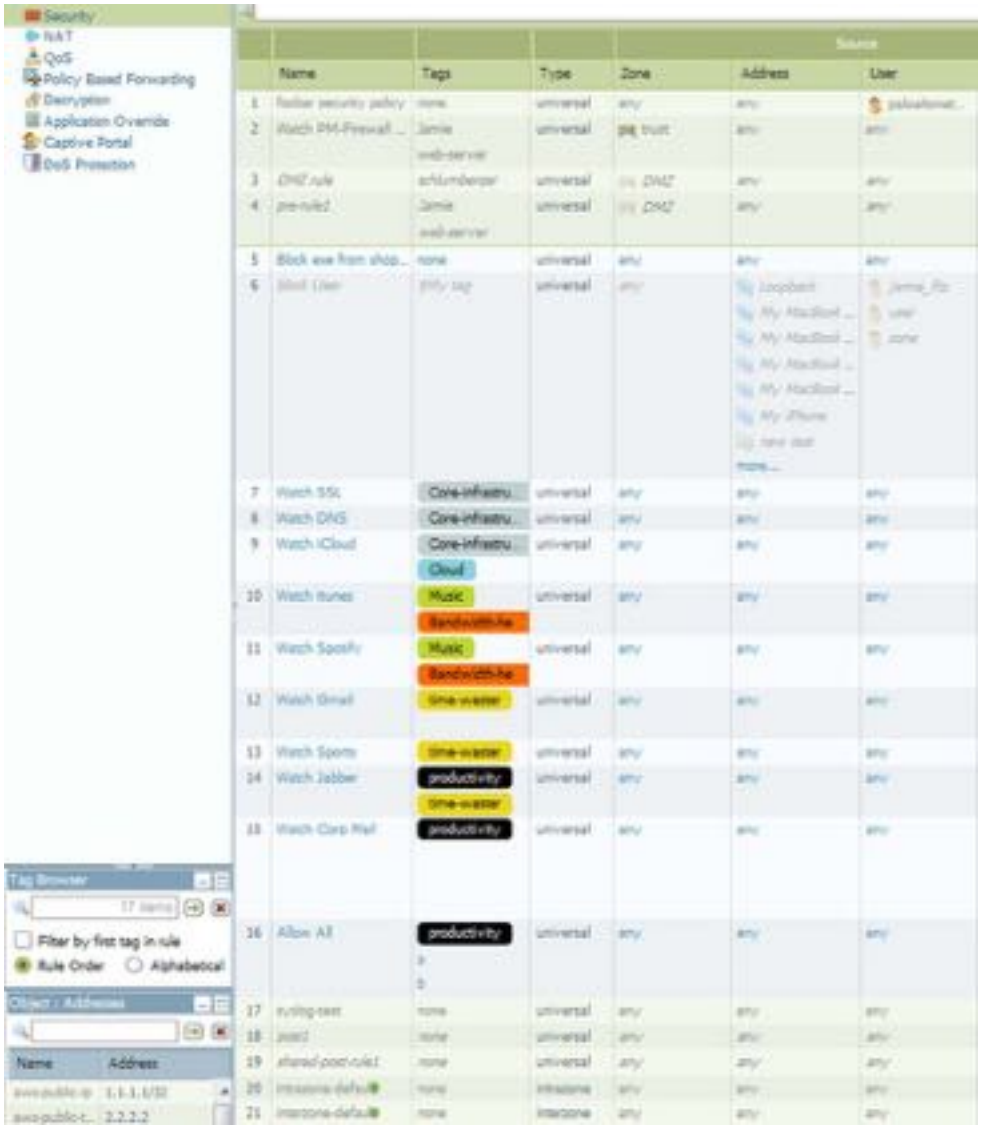

# Move or Clone a Policy Rule or Object to a Different Virtual System

On a firewall that has more than one virtual system (vsys), you can move or clone policy rules and objects to a different vsys or to the Shared location. Moving and cloning save you the effort of deleting, recreating, or renaming rules and objects. If the policy rule or object that you will move or clone from a vsys has references to objects in that vsys, move or clone the referenced objects also. If the references are to shared objects, you do not have to include those when moving or cloning. You can perform a Global Find to check for references.

| Move or Clone a Policy Rule or Object to a Virtual System | |
|---|---|
| Step 1 | Select the policy type (for example, **Policy > Security**) or object type (for example, **Objects > Addresses**). |
| Step 2 | Select the **Virtual System** and select one or more policy rules or objects. |
| Step 3 | Perform one of the following steps:<br>• Select **Move > Move to other vsys** (for policy rules).<br>• Click **Move** (for objects).<br>• Click **Clone** (for policy rules or objects). |
| Step 4 | In the **Destination** drop-down, select the new virtual system or **Shared**. The default is the **Virtual System** selected in Step 2. |
| Step 5 | (Policy rules only) Select the **Rule order**:<br>• **Move top** (default)—The rule will come before all other rules.<br>• **Move bottom**—The rule will come after all other rules.<br>• **Before rule**—In the adjacent drop-down, select the rule that comes after the Selected Rules.<br>• **After rule**—In the adjacent drop-down, select the rule that comes before the Selected Rules. |
| Step 6 | The **Error out on first detected error in validation** check box is selected by default. The firewall stops performing the checks for the move or clone action when it finds the first error, and displays just this error. For example, if an error occurs when the **Destination** vsys doesn't have an object that the policy rule you are moving references, the firewall will display the error and stop any further validation. When you move or clone multiple items at once, selecting this check box will allow you to find one error at a time and troubleshoot it. If you clear the check box, the firewall collects and displays a list of errors. If there are any errors in validation, the object is not moved or cloned until you fix all the errors. |
| Step 7 | Click **OK** to start the error validation. If the firewall displays errors, fix them and retry the move or clone operation. If the firewall doesn't find errors, the object is moved or cloned successfully. After the operation finishes, click **Commit**. |


© Palo Alto Networks, Inc.

# Use Tags to Group and Visually Distinguish Objects

You can tag objects to group related items and add color to the tag in order to visually distinguish tagged objects. Tags can be added to the following objects: address objects, address groups, zones, service groups, and policy rules.

The firewall and Panorama support both static tags and dynamic tags, dynamic tags are registered from a variety of sources and are not displayed with the static tags, because dynamic tags are not part of the device configuration. See Register IP Addresses and Tags Dynamically for information on registering tags dynamically. The tags discussed in this section are statically added and are part of the device configuration.

One or more tags can be applied to objects and to policy rules; a maximum of 64 tags can be applied to an object. Panorama supports a maximum of 10,000 tags that can be apportioned across Panorama (shared and device groups) and the managed devices (including devices with multiple virtual systems).

To use tags effectively, see the following topics:

- Create and Apply Tags
- Modify Tags
- Use the Tag Browser

# Create and Apply Tags

| Create and Apply tags | |
|---|---|
| Step 1 Create tags.<br><br>To tag a zone, you must create a tag with the same name as the zone. When the zone is attached in policy rules, the tag color automatically displays as the background color against the zone name. | 1. Select **Objects > Tags**.<br>2. On Panorama or a multiple virtual system firewall, select the **Device Group** or the **Virtual System** to which this object must belong.<br>3. Click **Add** and enter a **Name** to identify the tag. The maximum length is 127 characters.<br>4. (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.<br>5. (Optional) Assign one of the 16 predefined colors to the tag. By default, no color is selected.<br> 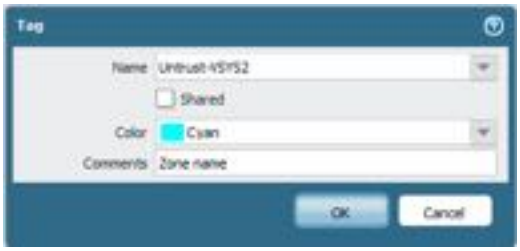 <br>6. Click **OK** and **Commit** to save the changes. |
| Step 2 Apply tags to policy. | 1. Select **Policies** and any rulebase under it.<br>2. Click **Add** to create a policy rule and use the tagged objects you created in Step 1.<br>3. Verify that the tags are in use.<br> |
| Step 3 Apply tags to an address object, address group, service, or service group. | 1. Create the object.<br>For example to create a service group, select **Objects > Service Groups > Add**.<br>2. Select the tag(s) from the **Tag** drop-down or enter a phrase to create a new tag.<br>To edit a tag or add color to the tag, see Modify Tags. |

 © Palo Alto Networks, Inc.

## Modify Tags

| Modify Tags |
| --- |
| • Select **Objects > Tags** to perform any of the following operations with tags:<br>  • Click the link in the **Name** column to edit the properties of a tag.<br>  • Select a tag in the table, and click **Delete** to remove the tag from the firewall.<br>  • Click **Clone** to create a duplicate tag with the same properties. A numerical suffix is added to the tag name. For example, FTP-1. |

For details on creating tags, see Create and Apply Tags. For information on working with tags, see Use the Tag Browser.

## Use the Tag Browser

The tag browser provides a way to view all the tags used within a rulebase. In rulebases with a large number of rules, the tag browser simplifies the display by presenting the tags, the color code, and the rule numbers in which the tags are used.

It also allows you to group rules using the first tag applied to the rule. As a best practice, use the first tag to identify the primary purpose for a rule. For example, the first tag can identify a rule by a high-level function such as best practice, or Internet access or IT sanctioned applications or high-risk applications. In the tag browser, when you **Filter by first tag in rule**, you can easily identify gaps in coverage and move rules or add new rules within the rulebase. All the changes are saved to the candidate configuration until you commit the changes on the firewall and make them a part of the running configuration.

For devices that are managed by Panorama, the tags applied to pre-rules and post-rules that have been pushed from Panorama, display in a green background and are demarcated with green lines so that you can identify these tags from the local tags on the device.

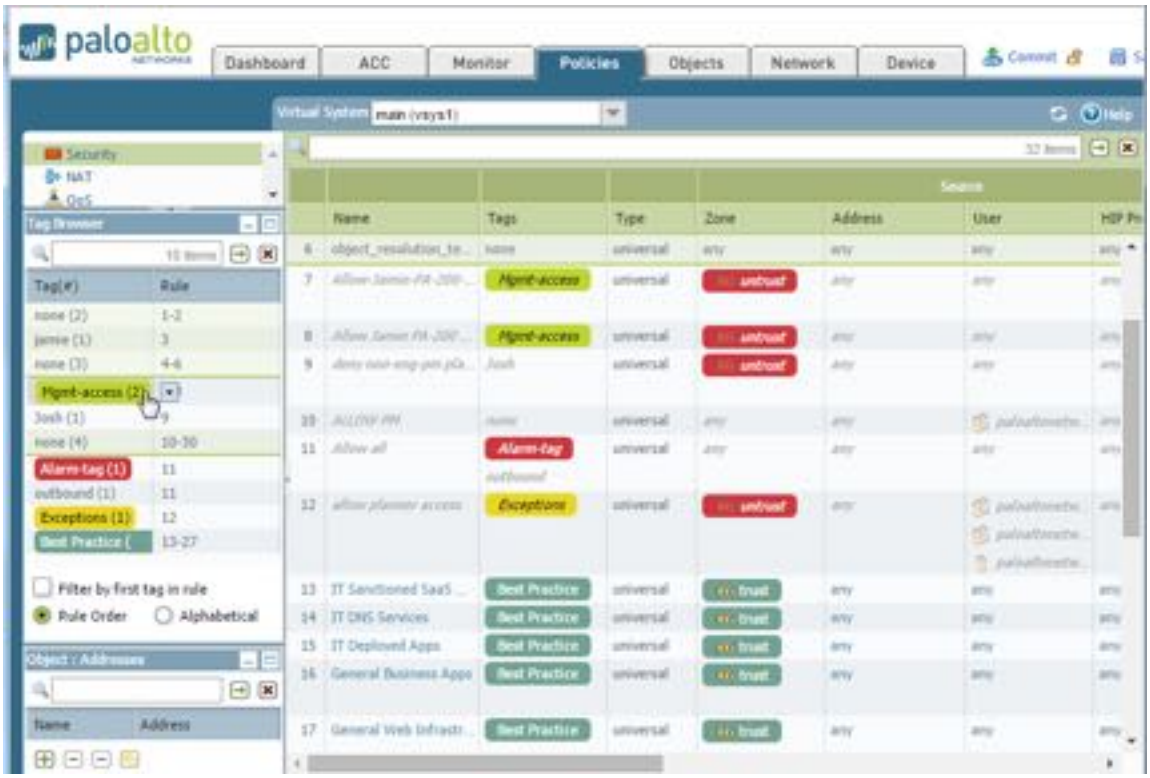


 © Palo Alto Networks, Inc.

| Use the Tag Browser | |
|---|---|
| • Explore the tag browser. | 1. Access the **Tag Browser** on the left pane of the **Policies >** tab. The tag browser displays the tags that have been used in the rules for the selected rulebase, for example **Policies > Security**.<br>2. **Tag (#)**—Displays the label and the rule number or range of numbers in which the tag is used contiguously. Hover over the label to see the location where the rule was defined, it can be inherited from a shared location, a device group, or a virtual system.<br>3. **Rule**—Lists the rule number or range of numbers associated with the tags.<br>4. Sort the tags.<br>• **Filter by first tag in rule**—Sorts rules using the first tag applied to each rule in the rulebase. This view is particularly useful if you want to narrow the list and view related rules that might be spread around the rulebase. For example if the first tag in each rule denotes its function—best practices, administration, web-access, data center access, proxy—you can narrow the result and scan the rules based on function.<br>• **Rule Order**—Sorts the tags in the order of appearance within the selected rulebase. When displayed in order of appearance, tags used in contiguous rules are grouped. The rule number with which the tag is associated is displayed along with the tag name.<br>• **Alphabetical**—Sorts the tags in alphabetical order within the selected rulebase. The display lists the tag name and color (if a color is assigned) and the number of times it is used within the rulebase.<br>The label **None** represents rules without any tags; it does not display rule numbers for untagged rules. When you select **None**, the right pane is filtered to display rules that have no tags assigned to them.<br>5. **Clear**—Clears the filter on the currently selected tags in the search bar.<br>6. **Search bar**—To search for a tag, enter the term and click the green arrow icon to apply the filter. It also displays the total number of tags in the rulebase and the number of selected tags.<br>7. Expand or collapse the tag browser. |

| Use the Tag Browser (Continued) | |
|---|---|
| • Tag a rule. | 1. Select a rule on the right pane.<br>2. Do one of the following:<br>• Select a tag in the tag browser and select **Apply the Tag to the Selection(s)** from the drop-down. 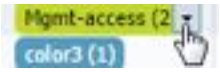 <br>• Drag and drop tag(s) from the tag browser on to the Tags column of the rule. When you drop a tag, a confirmation dialog displays.<br>3. **Commit** the changes. |
| • View rules that match the selected tags.<br>You can filter rules based on tags with an AND or an OR operator. | • OR filter: To view rules that have specific tags, select one or more tags in the tag browser; the right pane only displays the rules that include any of the currently selected tags.<br>• AND filter: To view rules that have all the selected tags, hover over the number associated with the tag in the **Rule** column of the tag browser and select **Filter**. Repeat to add more tags. <br>Click the apply filter icon in the search bar on the right pane. The results are displayed using an AND operator. |
| • View the currently selected tags. | To view the currently selected tags, hover over the **Clear** label in the tag browser. |
| • Untag a rule. | Hover over the rule number associated with a tag in the **Rule** column of the tag browser and select **Untag Rule(s)**. Confirm that you want to remove the selected tag from the rule. **Commit** the changes. |
| • Reorder rules using tags. | Select one or more tags and hover over the rule number in the Rule column of the tag browser and select **Move Rule(s)**.<br>Select a tag from the drop-down in the move rule window and select whether you want to **Move Before** or **Move After** the tag selected in the drop-down. **Commit** the changes. |

 © Palo Alto Networks, Inc.

| Use the Tag Browser (Continued) | |
|---|---|
| • Add a new rule that applies the selected tags. | Select one or more tags and hover over the rule number in the **Rule** column of the tag browser, and select **Add New Rule**. Define the rule and **Commit** the changes.<br>The numerical order of the new rule varies by whether you selected a rule on the right pane. If you did not select a rule on the right pane, the new rule will be added after the rule to which the selected tag(s) belongs. Otherwise, the new rule is added after the selected rule. |
| • Search for a tag. | In the tag browser, enter the first few letters of the tag name you want to search for and click the Apply Filter icon. The tags that match your input will display. |

# Use a Dynamic Block List in Policy

The firewall or Panorama typically enforce policy for a source or destination IP address that is defined as a static object on the firewall. If you need agility in enforcing policy for a list of source/destination IP addresses that emerge ad hoc, you can use dynamic block lists.

A dynamic block list is a text file that contains a list of IP addresses, IP ranges, or IP subnets, and is hosted on a web server. The dynamic block list can be used to deny or allow access to the IP addresses (IPv4 and IPv6) included in the list. For example, you can use it as a whitelist for allowing a set of IP addresses or as a blacklist to disallow access to the specified IP addresses. At a configured interval, the firewall dynamically imports the list and enforces policy for the IP addresses included in the list. When you modify the list, the firewall retrieves the updates; a configuration change or commit is not required on the firewall. If the web server is unreachable, the firewall or Panorama will use the last successfully retrieved list for enforcing policy until the connection is restored with the web server that hosts the list.


- View the IP Address Limit For Your Firewall Model
- Formatting Guidelines for Dynamic Block Lists
- Enforce Policy with a Dynamic Block List
- View the List of IP addresses in the Dynamic Block List
- Retrieve a Dynamic Block List from Web Server


## View the IP Address Limit For Your Firewall Model

Irrespective of the firewall model, each firewall supports a maximum of 10 Dynamic Block Lists.

To find the maximum number of addresses, address groups, and IP addresses per group, for your model of the firewall, use the following CLI command:

**show system state | match cfg.general.max-address**

For example:

```
admin@PA-7050> show system state | match cfg.general.max-address
cfg.general.max-address: 80000
cfg.general.max-address-group: 8000
cfg.general.max-address-per-group: 500
```

Each list can contain the maximum number of addresses supported by your firewall model minus 300. Up to 300 IP addresses are reserved for internal use on the firewall and are deducted from the available limit. Therefore, in the example above, the firewall can have a maximum of 79,700 IP addresses.

 © Palo Alto Networks, Inc.

## Formatting Guidelines for Dynamic Block Lists

The dynamic block list can include individual IP addresses, subnet addresses (address/mask), or range of IP addresses. In addition, the block list can include comments and special characters such as * , : , ; , #, or /. The syntax for each line in the list is `[IP address, IP/Mask, or IP start range-IP end range] [space] [comment].`

Because the firewall ignores incorrectly formatted lines, use these guidelines when defining the list:

- Enter each IP address/range/subnet in a new line; URLs are not supported in this list.
- If you add comments, the comment must be on the same line as the IP address/range/subnet. The space at the end of the IP address is the delimiter that separates a comment from the IP address.

An example:

```
192.168.20.10/32
2001:db8:123:1::1 #test IPv6 address
192.168.20.0/24 ; test internal subnet
2001:db8:123:1::/64 test internal IPv6 range
192.168.20.40-192.168.20.50
```

For an IP address that is blocked, you can display a notification page only if the protocol is HTTP.

## Enforce Policy with a Dynamic Block List

| Enforce Policy with a Dynamic Block List | | |
|---|---|---|
| Step 1 | Create the dynamic block list and host it on a web server, so that the firewall can retrieve the list for policy evaluation. | 1. Create a text file and enter the IP addresses for which you want to enforce policy. For syntax, see Formatting Guidelines for Dynamic Block Lists. |
| Step 2 | Create a dynamic block list object on the firewall. | 1. Select **Objects > Dynamic Block Lists**.<br>2. Click **Add** and enter a descriptive **Name** for the list.<br>3. (Optional) Select **Shared**, to share the list with all virtual systems on a device that is enabled for multiple virtual systems. By default, the object is created on the virtual system that is currently selected in the **Virtual Systems** drop-down.<br>4. Enter the **Source** URL (hostname or IP address and the path) for the list you just created on the web server. For example, https://1.2.3.4/DBL_2014<br>5. Click **Test Source URL** to verify that the firewall or Panorama can connect to the web server.<br>6. (Optional) Specify the **Repeat** frequency at which the firewall or Panorama must retrieve the list. By default the list is retrieved ever hour.<br>7. Click **OK** to save the changes. |

| Enforce Policy with a Dynamic Block List | | |
|---|---|---|
| Step 3 | Use the dynamic block list as a source or destination address object in policy.<br> Create separate dynamic block lists if you want to specify allow and deny actions for specific IP addresses.<br>The list can be referenced in any policy type. In this example, we attach it as a destination object in security policy. | 1. Select **Policies > Security**.<br>2. Click **Add** and give the rule a descriptive name in the **General** tab.<br>3. In the **Source** tab, select the **Source Zone**.<br>4. In the **Destination** tab, select the **Destination Zone** and select the dynamic block list as the Destination Address.<br>5. In the **Service/ URL Category** tab, make sure the **Service** is set to **application-default**.<br>6. In the **Actions** tab, set the **Action Setting** to Allow or Deny.<br>7. Leave all the other options at the default values.<br>8. Click **OK** to save the changes.<br>9. **Commit** the changes. |
| Step 4 | Test that the policy action is enforced. | 1. Access a IP address that is included in the dynamic block list and verify that action you defined is enforced.<br>2. Select **Monitor > Logs > Traffic** and see the log entry for the session.<br>3. To verify the policy rule that matches a flow, use the following CLI command:<br>**test security-policy-match source** *<IP_address>* **destination** *<IP_address>* **destination port** *<port_number>* **protocol** *<protocol_number>* |

## View the List of IP addresses in the Dynamic Block List

**View the IP Addresses Included in the Dynamic Block List**

To view the list of IP addresses that the firewall has retrieved from the web server enter the following CLI command:

**request system external-list show name *<name>***

For example, for a list named case DBL_2014, the output is:

```
vsys1/DBL_2014:
  Next update at: Wed Aug 27 16:00:00 2014
  IPs:
  1.1.1.1
  1.2.2.2/20 #test China
  192.168.255.0; test internal
  192.168.254.0/24 test internal range
```

 © Palo Alto Networks, Inc.

## Retrieve a Dynamic Block List from Web Server

The firewall or Panorama can be configured to retrieve the list from the web server on an hourly, daily, weekly, or monthly basis. If you have added or deleted IP addresses on the list and need to trigger an immediate refresh, you must use the Command Line Interface.

**Retrieve a Dynamic Block List**

1. Enter the command: **request system external-list refresh name *<name>***

   For example, `request system external-list refresh name DBL_2014`
2. Get the job ID for the refresh job using the CLI command: **show jobs all**

   Look for the last EBL Refresh job in the list.
3. View the details for the job ID. Use the command **show jobs id *<number>***

   A message indicating the success or failure displays. For example:

```
admin@PA-200> show jobs id 55
Enqueued                    ID             Type    Status Result Completed
------------------------------------------------------------------------------
2014/08/26 15:34:14         55       EBLRefresh       FIN     OK 15:34:40
Warnings:
Details:
```

# Register IP Addresses and Tags Dynamically

To mitigate the challenges of scale, lack of flexibility and performance, the architecture in networks today allows for clients, servers, and applications to be provisioned, changed, and deleted on demand. This agility poses a challenge for security administrators because they have limited visibility into the IP addresses of the dynamically provisioned clients and servers, and the plethora of applications that can be enabled on these virtual resources.

The firewall (hardware-based platforms and the VM-Series) supports the ability to register IP addresses and tags dynamically. The IP addresses and tags can be registered on the firewall directly or registered on the firewall through Panorama.This dynamic registration process can be enabled using any of the following options:

- **User-ID agent for Windows**—In an environment where you've deployed the User-ID agent, you can enable the User-ID agent to monitor up to 100 VMware ESXi and/or vCenter Servers. As you provision or modify virtual machines on these VMware servers, the agent can retrieve the IP address changes and share them with the firewall.
- **VM Information Sources**—Allows you to monitor VMware ESXi and vCenter Server, and the AWS-VPC to retrieve IP address changes when you provision or modify virtual machines on these sources. VM Information Sources polls for a predefined set of attributes and does not require external scripts to register the IP addresses through the XML API. See Monitor Changes in the Virtual Environment.
- **VMware Service Manager** (only available for the integrated NSX solution)—The integrated NSX solution is designed for automated provisioning and distribution of Palo Alto Networks next-generation security services and the delivery of dynamic context-based security policies using Panorama. The NSX Manager updates Panorama with the latest information on the IP addresses and tags associated with the virtual machines deployed in this integrated solution. For information on this solution, see Set Up a VM-Series NSX Edition Firewall.
- **XML API**—The firewall and Panorama support an XML API that uses standard HTTP requests to send and receive data. You can use this API to register IP addresses and tags with the firewall or Panorama. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports REST-based services. Refer to the PAN-OS XML API Usage Guide for details.

For information on creating and using Dynamic Address Groups, see Use Dynamic Address Groups in Policy.

For the CLI commands for registering tags dynamically, see CLI Commands for Dynamic IP Addresses and Tags.

 © Palo Alto Networks, Inc.

# Monitor Changes in the Virtual Environment

To secure applications and prevent threats in an environment where new users and servers are constantly emerging, your security policy must be nimble. To be nimble, the firewall must be able to learn about new or modified IP addresses and consistently apply policy without requiring configuration changes on the firewall.

This capability is provided by the coordination between the **VM Information Sources** and **Dynamic Address Groups** features on the firewall. The firewall and Panorama provide an automated way to gather information on the virtual machine (or guest) inventory on each monitored source and create policy objects that stay in sync with the dynamic changes on the network.


- Enable VM Monitoring to Track Changes on the Virtual Network
- Attributes Monitored in the AWS and VMware Environments
- Use Dynamic Address Groups in Policy

# Enable VM Monitoring to Track Changes on the Virtual Network

VM information sources provides an automated way to gather information on the Virtual Machine (VM) inventory on each monitored source (host); the firewall can monitor the VMware ESXi and vCenter Server, and the AWS-VPC. As virtual machines (guests) are deployed or moved, the firewall collects a predefined set of attributes (or metadata elements) as tags; these tags can then be used to define Dynamic Address Groups (see Use Dynamic Address Groups in Policy) and matched against in policy.

Up to 10 VM information sources can be configured on the firewall or pushed using Panorama templates. By default, the traffic between the firewall and the monitored sources uses the management (MGT) port on the firewall.

**VM Information Sources** offers easy configuration and enables you to monitor a predefined set of 16 metadata elements or attributes. See Attributes Monitored in the AWS and VMware Environmentsfor the list.

 © Palo Alto Networks, Inc.

**Set up the VM Monitoring Agent**

Step 1 Enable the VM Monitoring Agent.

Up to 10 sources can be configured for each firewall, or for each virtual system on a multiple virtual systems capable firewall.

If your firewalls are configured in a high availability configuration:

- An active/passive setup, only the active firewall monitors the VM sources.
- An active/active setup, only the firewall with the priority value of primary monitors the VM sources.

1. Select **Device > VM Information Sources**.
2. Click **Add** and enter the following information:
   - A **Name** to identify the VMware ESX(i) or vCenter Server that you want to monitor.
   - Enter the **Host information for the server—**hostname or IP address and the **Port** on which it is listening.
   - Select the **Type** to indicate whether the source is a **VMware ESX(i)** server or a **VMware vCenter** Server.
   - Add the credentials (**Username** and **Password**) to authenticate to the server specified above.
   - Use the credentials of an administrative user to enable access.
   - (Optional) Modify the **Update interval** to a value between 5-600 seconds. By default, the firewall polls every 5 seconds. The API calls are queued and retrieved within every 60 seconds, so updates may take up to 60 seconds plus the configured polling interval.

   - (Optional) Enter the interval in hours when the connection to the monitored source is closed, if the host does not respond. (default: 2 hours, range 2-10 hours)
     To change the default value, select the check box to **Enable timeout when the source is disconnected** and specify the value. When the specified limit is reached or if the host cannot be accessed or does not respond, the firewall will close the connection to the source.
   - Click **OK**, and **Commit** the changes.
   - Verify that the connection **Status** displays as connected .

| Set up the VM Monitoring Agent (Continued) | |
|---|---|
| Step 2 Verify the connection status. | Verify that the connection **Status** displays as connected.<br> <br>If the connection status is pending or disconnected, verify that the source is operational and that the firewall is able to access the source. If you use a port other than the MGT port for communicating with the monitored source, you must change the service route (**Device > Setup > Services**, click the **Service Route Configuration** link and modify the **Source Interface** for the **VM Monitor** service). |

 © Palo Alto Networks, Inc.

## Attributes Monitored in the AWS and VMware Environments

Each VM on a monitored ESXi or vCenter server must have VMware Tools installed and running. VMware Tools provide the capability to glean the IP address(es) and other values assigned to each VM.

In order to collect the values assigned to the monitored VMs, the firewall monitors the following predefined set of attributes:

| Attributes Monitored on a VMware Source | Attributes Monitored on the AWS-VPC |
|---|---|
| • UUID | • Architecture |
| • Name | • Guest OS |
| • Guest OS | • Image ID |
| • VM State — the power state can be poweredOff, poweredOn, standBy, and unknown. | • Instance ID |
| • Annotation | • Instance State |
| • Version | • Instance Type |
| • Network — Virtual Switch Name, Port Group Name, and VLAN ID | • Key Name |
| • Container Name —vCenter Name, Data Center Object Name, Resource Pool Name, Cluster Name, Host, Host IP address. | • Placement—Tenancy, Group Name, Availability Zone |
| | • Private DNS Name |
| | • Public DNS Name |
| | • Subnet ID |
| | • Tag (key, value) (up to5 tags supported per instance |
| | • VPC ID |

## Use Dynamic Address Groups in Policy

Dynamic address groups are used in policy. They allow you to create policy that automatically adapts to changes—adds, moves, or deletions of servers. It also enables the flexibility to apply different rules to the same server based on *tags* that define its role on the network, the operating system, or the different kinds of traffic it processes.

A dynamic address group uses tags as a filtering criteria to determine its members. The filter uses logical *and* and *or* operators. All IP addresses or address groups that match the filtering criteria become members of the dynamic address group. Tags can be defined statically on the firewall and/or registered (dynamically) to the firewall. The difference between static and dynamic tags is that static tags are part of the configuration on the firewall, and dynamic tags are part of the runtime configuration. This implies that a commit is not required to update dynamic tags; the tags must however be used by Dynamic Address Groups that are referenced in policy, and the policy must be committed on the device.

To dynamically register tags, you can use the XML API or the VM Monitoring agent on the firewall or on the User-ID agent. Each tag is a metadata element or attribute-value pair that is registered on the firewall or Panorama. For example, IP1 {tag1, tag2,.....tag32}, where the IP address and the associated tags are maintained as a list; each registered IP address can have up to 32 tags such as the operating system, the datacenter or the virtual switch to which it belongs. Within 60 seconds of the API call, the firewall registers the IP address and associated tags, and automatically updates the membership information for the dynamic address group(s).

The maximum number of IP addresses that can be registered for each platform is different. Use the following table for specifics on your platform:

| Platform | Maximum number of dynamically registered IP addresses |
|---|---|
| PA-7000 Series, PA-5060, VM-1000-HV | 100,000 |
| PA-5050 | 50,000 |
| PA-5020 | 25,000 |
| PA-4000 Series, PA-3000 Series | 5,000 |
| PA-2000 Series, PA-500, PA-200, VM-300, VM-200, VM-100 | 1,000 |

The following example shows how dynamic address groups can simplify network security enforcement. The example workflow shows how to:

- Enable the VM Monitoring agent on the firewall, to monitor the VMware ESX(i) host or vCenter Server and register VM IP addresses and the associated tags.
- Create dynamic address groups and define the tags to filter. In this example, two address groups are created. One that only filters for dynamic tags and another that filters for both static and dynamic tags to populate the members of the group.
- Validate that the members of the dynamic address group are populated on the firewall.
- Use dynamic address groups in policy. This example uses two different security policies:
  - A security policy for all Linux servers that are deployed as FTP servers; this rule matches on dynamically registered tags.

 © Palo Alto Networks, Inc.

– A security policy for all Linux servers that are deployed as web servers; this rule matches on a dynamic address group that uses static and dynamic tags.

- Validate that the members of the dynamic address groups are updated as new FTP or web servers are deployed. This ensure that the security rules are enforced on these new virtual machines too.

| **Use Dynamic Address Groups in Policy** | |
|---|---|
| Step 1 Enable VM Source Monitoring. | See Enable VM Monitoring to Track Changes on the Virtual Network. |
| Step 2 Create dynamic address groups on the firewall.<br><br>View the tutorial to see a big picture view of the feature. | 1. Log in to the web interface of the firewall.<br>2. Select **Object > Address Groups**.<br>3. Click **Add** and enter a **Name** and a **Description** for the address group.<br>4. Select **Type** as **Dynamic**.<br>5. Define the match criteria. You can select dynamic and static tags as the match criteria to populate the members of the group. Click **Add Match Criteria**, and select the **And** or **Or** operator and select the attributes that you would like to filter for or match against. and then click **OK**.<br> <br>6. Click **Commit**. |

The match criteria for each dynamic address group in this example is as follows:

ftp_server: matches on the guest operating system “Linux 64-bit” and annotated as “ftp” ('guestos.Ubuntu Linux 64-bit' and 'annotation.ftp').

web-servers: matches on two criteria—the tag black or if the guest operating system is Linux 64-bit and the name of the server us Web_server_Corp. ('guestos.Ubuntu Linux 64-bit' and 'vmname.WebServer_Corp' or 'black')

**Use Dynamic Address Groups in Policy (Continued)**

Step 3 Use dynamic address groups in policy.

View the tutorial.

1. Select **Policies > Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the Dynamic address group you created in Step 2 above.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeats Steps 1 through 6 above to create another policy rule.
8. Click **Commit**.

This example shows how to create two policies: one for all access to FTP servers and the other for access to web servers.

Step 4 Validate that the members of the dynamic address group are populated on the firewall.

1. Select **Policies > Security**, and select the rule.
2. Select the drop-down arrow next to the address group link, and select **Inspect**. You can also verify that the match criteria is accurate.

3. Click the **more** link and verify that the list of registered IP addresses is displayed.

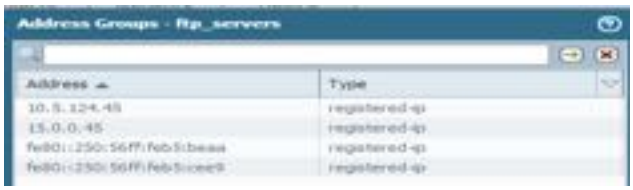


Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.


© Palo Alto Networks, Inc.

# CLI Commands for Dynamic IP Addresses and Tags

The Command Line Interface on the firewall and Panorama give you a detailed view into the different sources from which tags and IP addresses are dynamically registered. It also allows you to audit registered and unregistered tags. The following examples illustrate the capabilities in the CLI.

| Example | CLI Command |
|---|---|
| View all registered IP addresses that match the tag, `state.poweredOn` or that are not tagged as `vSwitch0` | `show log iptag tag_name equal state.poweredOn`<br>`show log iptag tag_name not-equal switch.vSwitch0` |
| View all dynamically registered IP addresses that were sourced by VM Information Source with name `vmware1` and tagged as `poweredOn` | `show vm-monitor source source-name vmware1 tag state.poweredOn registered-ip all`<br><br>`registered IP Tags`<br>`---------- ----------`<br>`fe80::20c:29ff:fe69:2f76 "state.poweredOn"`<br>`10.1.22.100 "state.poweredOn"`<br>`2001:1890:12f2:11:20c:29ff:fe69:2f76 "state.poweredOn"`<br>`fe80::20c:29ff:fe69:2f80 "state.poweredOn"`<br>`192.168.1.102 "state.poweredOn"`<br>`10.1.22.105 "state.poweredOn"`<br>`2001:1890:12f2:11:2cf8:77a9:5435:c0d "state.poweredOn"`<br>`fe80::2cf8:77a9:5435:c0d "state.poweredOn"` |
| Clear all IP addresses and tags learned from a specific VM Monitoring source without disconnecting the source. | `debug vm-monitor clear source-name <name>` |
| Display IP addresses registered from all sources. | `show object registered-ip all` |
| Display the count for IP addresses registered from all sources. | `show object registered-ip all option count` |
| Clear IP addresses registered from all sources | `debug object registered-ip clear all` |
| Add or delete tags for a given IP address that was registered using the XML API. | `debug object test registered-ip [<register/unregister>] <ip/netmask> <tag>` |

| Example | CLI Command |
|---|---|
| View all tags registered from a specific information source. | `show vm-monitor source source-name vmware1 tag all`<br>`vlanId.4095`<br>`vswitch.vSwitch1`<br>`host-ip.10.1.5.22`<br>`portgroup.TOBEUSED`<br>`hostname.panserver22`<br>`portgroup.VM Network 2`<br>`datacenter.ha-datacenter`<br>`vlanId.0`<br>`state.poweredOn`<br>`vswitch.vSwitch0`<br>`vmname.Ubuntu22-100`<br>`vmname.win2k8-22-105`<br>`resource-pool.Resources`<br>`vswitch.vSwitch2`<br>`guestos.Ubuntu Linux 32-bit`<br>`guestos.Microsoft Windows Server 2008 32-bit`<br>`annotation.`<br>`version.vmx-08`<br>`portgroup.VM Network`<br>`vm-info-source.vmware1`<br>`uuid.564d362c-11cd-b27f-271f-c361604dfad7`<br>`uuid.564dd337-677a-eb8d-47db-293bd6692f76`<br>`Total: 22` |
| View all tags registered from a specific data source, for example from the VM Monitoring Agent on the firewall, the XML API, Windows User-ID Agent or the CLI. | • To view tags registered from the CLI:<br>`show log iptag datasource_type equal unknown`<br>• To view tags registered from the XML API:<br>`show log iptag datasource_type equal xml-api`<br>• To view tags registered from VM Information sources:<br>`show log iptag datasource_type equal vm-monitor`<br>• To view tags registered from the Windows User-ID agent:<br>`show log iptag datasource_type equal xml-api datasource_subtype equal user-id-agent` |
| View all tags that are registered for a specific IP address (across all sources). | `debug object registered-ip show tag-source ip` *`ip_address`* `tag all` |

 © Palo Alto Networks, Inc.

# Identify Users Connected through a Proxy Server

If you have a proxy server deployed between the users on your network and the firewall, in HTTP/HTTPS requests the firewall might see the proxy server IP address as the source IP address in the traffic that the proxy forwards rather than the IP address of the client that requested the content. In many cases, the proxy server adds an X-Forwarded-For (XFF) header to traffic packets that includes the actual IPv4 or IPv6 address of the client that requested the content or from whom the request originated. In such cases, you can configure the firewall to read the XFF header values and determine the IP addresses of the client who requested the content. The firewall matches the XFF IP addresses with usernames that your policy rules reference so that those rules can control access for the associated users and groups. The firewall also uses the XFF-derived usernames to populate the source user fields of logs so you can monitor user access to web services.

You can also configure the firewall to add XFF values to URL Filtering logs. In these logs, an XFF value can be the client IP address, client username (if available), the IP address of the last proxy server traversed in a proxy chain, or any string of up to 128 characters that the XFF header stores.

XFF user identification applies only to HTTP or HTTPS traffic, and only if the proxy server supports the XFF header. If the header has an invalid IP address, the firewall uses that IP address as a username for group mapping references in policies. If the XFF header has multiple IP addresses, the firewall uses the first entry from the left.

- Use XFF Values for Policies and Logging Source Users
- Add XFF Values to URL Filtering Logs

## Use XFF Values for Policies and Logging Source Users

You can configure the firewall to use XFF values in user-based policies and in the source user fields of logs. To use XFF values in policies, you must also Map IP Addresses to Users, Map Users to Groups (if you have group-based policies), and configure policies based on users or groups.

Logging XFF values doesn't populate the source IP address values of logs. When you view the logs, the source field displays the IP address of the proxy server if one is deployed between the user clients and the firewall. However, you can configure the firewall to Add XFF Values to URL Filtering Logs so that you can see user IP addresses in those logs.

To ensure that attackers can't read and exploit the XFF values in web request packets that exit the firewall to retrieve content from an external server, you can also configure the firewall to strip the XFF values from outgoing packets.

These options are not mutually exclusive: if you configure both, the firewall zeroes out XFF values only after using them in policies and logs.

| **Use XFF Values for Policies and Logging Source Users** | | |
|---|---|---|
| Step 1 | Enable the firewall to use XFF values in policies and in the source user fields of logs. | 1. Select **Device > Setup > Content-ID** and edit the X-Forwarded-For Headers settings.<br>2. Select the **Use X-Forwarded-For Header in User-ID** check box. |
| Step 2 | Remove XFF values from outgoing web requests. | 1. Select the **Strip X-Forwarded-For Header** check box.<br>2. Click **OK** and **Commit**. |

| Use XFF Values for Policies and Logging Source Users (Continued) | | |
|---|---|---|
| Step 3 | Verify the firewall is populating the source user fields of logs. | 1. Select a log type that has a source user field (for example, **Monitor > Logs > Traffic**).<br>2. Verify that the Source User column displays the usernames of users who access the web. |

## Add XFF Values to URL Filtering Logs

You can configure the firewall to add the XFF values from web requests to URL Filtering logs. The XFF values that the logs display can be client IP addresses, usernames if available, or any values of up to 128 characters that the XFF fields store.

This method of logging XFF values doesn't add usernames to the source user fields in URL Filtering logs. To populate the source user fields, see Use XFF Values for Policies and Logging Source Users.

| Add XFF Values to URL Filtering Logs | | |
|---|---|---|
| Step 1 | Configure a URL Filtering profile. | 1. Select **Objects > Security Profiles > URL Filtering**.<br>2. Select an existing profile or **Add** a new profile and enter a descriptive **Name**.<br>You can't enable XFF logging in the default URL Filtering profile.<br>3. In the **Categories** tab, Define how to control access to web content.<br>4. Select the **Settings** tab and select the **X-Forwarded-For** check box.<br>5. Click **OK** to save the profile. |
| Step 2 | Attach the URL Filtering profile to a policy rule. | 1. Select **Policies > Security** and click the rule.<br>2. Select the **Actions** tab, set the **Profile Type** to **Profiles**, and select the **URL Filtering** profile you just created.<br>3. Click **OK** and **Commit**. |
| Step 3 | Verify the firewall is logging XFF values. | 1. Select **Monitor > Logs > URL Filtering**.<br>2. Display the XFF values in one of the following ways:<br>• To display the XFF value for a single log—Click the icon for the log to displays its details. The HTTP Headers section displays the X-Forwarded-For value.<br>• To display the XFF values for all logs—Open the drop-down in any column header, select **Columns**, and select the **X-Forwarded-For** check box. The page then displays an X-Forwarded-For column. |

 © Palo Alto Networks, Inc.

# Policy-Based Forwarding

Normally, the firewall uses the destination IP address in a packet to determine the outgoing interface. The firewall uses the routing table associated with the virtual router to which the interface is connected to perform the route lookup. Policy-Based Forwarding (PBF) allows you to override the routing table, and specify the outgoing or *egress* interface based on specific parameters such as source or destination IP address, or type of traffic.


- PBF
- Create a Policy-Based Forwarding Rule
- Use Case: PBF for Outbound Access with Dual ISPs
- Use Case: PBF for Routing Traffic Through Virtual Systems

# PBF

PBF rules allow traffic to take an alternative path from the next hop specified in the route table, and are typically used to specify an egress interface for security or performance reasons. Let's say your company has two links between the corporate office and the branch office: a cheaper Internet link and a more expensive leased line. The leased line is a high-bandwidth, low-latency link. For enhanced security, you can use PBF to send applications that aren't encrypted traffic, such as FTP traffic, over the private leased line and all other traffic over the Internet link. Or, for performance, you can choose to route business-critical applications over the leased line while sending all other traffic, such as web browsing, over the cheaper link.

## Egress Path and Symmetric Return

Using PBF, you can direct traffic to a specific interface on the firewall, drop the traffic, or direct traffic to another virtual system (on systems enabled for multiple virtual systems).

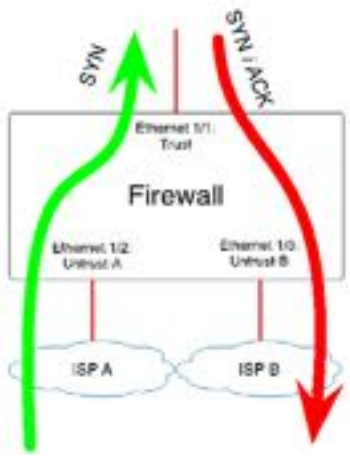


In networks with asymmetric routes, such as in a dual ISP environment, connectivity issues occur when traffic arrives at one interface on the firewall and leaves from another interface. If the route is asymmetrical, where the forward (SYN packet) and return (SYN/ACK) paths are different, the firewall is unable to track the state of the entire session and this causes a connection failure. To ensure that the traffic uses a symmetrical path, which means that the traffic arrives at and leaves from the same interface on which the session was created, you can enable the *Symmetric Return* option.

With symmetric return, the virtual router overrides a routing lookup for return traffic and instead directs the flow back to the MAC address from which it received the SYN packet (or first packet). However, if the destination IP address is on the same subnet as the ingress/egress interface's IP address, a route lookup is performed and symmetric return is not enforced. This behavior prevents traffic from being blackholed.

To determine the next hop for symmetric returns, the firewall uses an Address Resolution Protocol (ARP) table. The maximum number of entries that this ARP table supports is limited by the firewall model and the value is not user configurable. To determine the limit for your model, use the CLI command: **`show pbf return-mac all`**.

## Path Monitoring

Path monitoring allows you to verify connectivity to an IP address so that the firewall can direct traffic through an alternate route, when needed. The firewall uses ICMP pings as *heartbeats* to verify that the specified IP address is reachable.

A monitoring profile allows you to specify the threshold number of heartbeats to determine whether the IP address is reachable. When the monitored IP address is unreachable, you can either disable the PBF rule or specify a *fail-over* or *wait-recover* action. Disabling the PBF rule allows the virtual router to take over the routing

 © Palo Alto Networks, Inc.

decisions. When the fail-over or wait-recover action is taken, the monitoring profile continues to monitor whether the target IP address is reachable, and when it comes back up, the firewall reverts back to using the original route.

The following table lists the difference in behavior for a path monitoring failure on a new session versus an established session.

| Behavior of a session on a monitoring failure | If the rule stays enabled when the monitored IP address is unreachable<br>☐ Disable this rule if nexthop/monitor ip is unreachable | If rule is disabled when the monitored IP address is unreachable<br>☑ Disable this rule if nexthop/monitor ip is unreachable |
|---|---|---|
| For an established session | **wait-recover**—Continue to use egress interface specified in the PBF rule | **wait-recover**—Continue to use egress interface specified in the PBF rule |
| | **fail-over**—Use path determined by routing table (no PBF) | **fail-over**—Use path determined by routing table (no PBF) |
| For a new session | **wait-recover**—Use path determined by routing table (no PBF) | **wait-recover**—Check the remaining PBF rules. If no match, use the routing table |
| | **fail-over**—Use path determined by routing table (no PBF) | **fail-over**—Check the remaining PBF rules. If no match, use the routing table |

## Service Versus Applications in PBF

PBF rules are applied either on the first packet (SYN) or the first response to the first packet (SYN/ACK). This means that a PBF rule may be applied before the firewall has enough information to determine the application. Therefore, application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application.

However, if you specify an application in a PBF rule, the firewall performs *App-ID caching*. When an application passes through the firewall for the first time, the firewall does not have enough information to identify the application and therefore cannot enforce the PBF rule. As more packets arrive, the firewall determines the application and creates an entry in the App-ID cache and retains this App-ID for the session.When a new session is created with the same destination IP address, destination port, and protocol ID, the firewall could identify the application as the same from the initial session (based on the App-ID cache) and apply the PBF rule. Therefore, a session that is not an exact match and is not the same application, can be forwarded based on the PBF rule.

Further, applications have dependencies and the identity of the application can change as the firewall receives more packets. Because PBF makes a routing decision at the start of a session, the firewall cannot enforce a change in application identity. YouTube, for example, starts as web-browsing but changes to Flash, RTSP, or YouTube based on the different links and videos included on the page. However with PBF, because the firewall identifies the application as web-browsing at the start of the session, the change in application is not recognized thereafter.

PBF rules cannot be based on domain names; only IP addresses are valid; also, you cannot use custom-applications, application-filters or application groups in PBF rules.

# Create a Policy-Based Forwarding Rule

Use a PBF rule to direct traffic to a specific egress interface on the firewall, and override the default path for the traffic.

| Create a PBF Rule | |
|---|---|
| Step 1 Create a PBF rule.<br><br>When creating a PBF rule you must specify a name for the rule, a source zone or interface, and an egress interface. All other components are either optional or have a default value provided. | 1. Select **Policies > Policy Based Forwarding** and click **Add.**<br>2. Give the rule a descriptive name in the **General** tab.<br>3. In the **Source** tab, select the following:<br>a. Select the **Type**—**Zone** or **Interface**— to which the forwarding policy will be applied, and the relevant zone or interface.<br><br>PBF is only supported on Layer 3 interfaces.<br><br>b. (Optional) Specify the **Source Address** to which PBF will apply. For example, a specific IP address or subnet IP address from which you want to forward traffic to the interface or zone specified in this rule.<br><br>Use the **Negate** option to exclude a one or more source IP addresses from the PBF rule. For example, if your PBF rule directs all traffic from the specified zone to the Internet, **Negate** allows you to exclude internal IP addresses from the PBF rule.<br><br>The evaluation order is top down. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated.<br><br>c. (Optional) **Add** and select the **Source User** or groups of users to whom the policy applies.<br>4. In the **Destination/Application/Service** tab, select the following:<br>a. **Destination Address**. By default the rule applies to **Any** IP address. Use the **Negate** option to exclude one or more destination IP addresses from the PBF rule.<br><br>b. Select the Application(s) or Service(s) that you want to control using PBF.<br><br>Application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application. For more details, see Service Versus Applications in PBF. |


© Palo Alto Networks, Inc.

| Create a PBF Rule | |
|---|---|
| | 5. In the **Forwarding** tab, select the following:<br>a. Set the **Action.** The options are as follows:<br>– **Forward**—Directs the packet to a specific **Egress Interface**. Enter the **Next Hop** IP address for the packet.<br>– **Forward To VSYS**—(On a device enabled for multiple virtual systems) Select the virtual system to which to forward the packet.<br>– **Discard**—Drop the packet.<br>– **No PBF**—Exclude the packets that match the criteria for source/destination/application/service defined in the rule. Matching packets use the route table instead of PBF; the firewall uses the route table to exclude the matched traffic from the redirected port.<br>To trigger the specified action at a daily, weekly or non-recurring frequency, create and attach a **Schedule**.<br>b. (Optional) Enable Monitoring to verify connectivity to a target IP address or to the next hop IP address. Select **Monitor** and attach a monitoring **Profile** (default or custom) that specifies the action when the IP address is unreachable.<br>c. (Optional, required for asymmetric routing environments) Select **Enforce Symmetric Return** and enter one or more IP addresses in the **Next Hop Address List**.<br>Enabling symmetric return ensures that return traffic (say, from the Trust zone on the LAN to the Internet) is forwarded out through the same interface through which traffic ingresses from the Internet. |
| Step 2 Save the policies to the running configuration on the device. | Click **Commit**.<br>The PBF rule is in effect. |

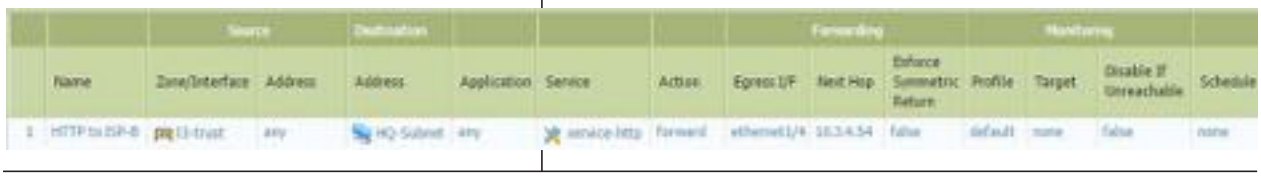

## Use Case: PBF for Outbound Access with Dual ISPs

In this use case, the branch office has a dual ISP configuration and implements PBF for redundant Internet access. The backup ISP is the default route for traffic from the client to the web servers. In order to enable redundant Internet access without using an internetwork protocol such as BGP, we use PBF with destination interface-based source NAT and static routes, and configure the firewall as follows:

- Enable a PBF rule that routes traffic through the primary ISP, and attach a monitoring profile to the rule. The monitoring profile triggers the firewall to use the default route through the backup ISP when the primary ISP is unavailable.
- Define Source NAT rules for both the primary and backup ISP that instruct the firewall to use the source IP address associated with the egress interface for the corresponding ISP. This ensures that the outbound traffic has the correct source IP address.
- Add a static route to the backup ISP, so that when the primary ISP is unavailable, the default route comes into effect and the traffic is directed through the backup ISP.

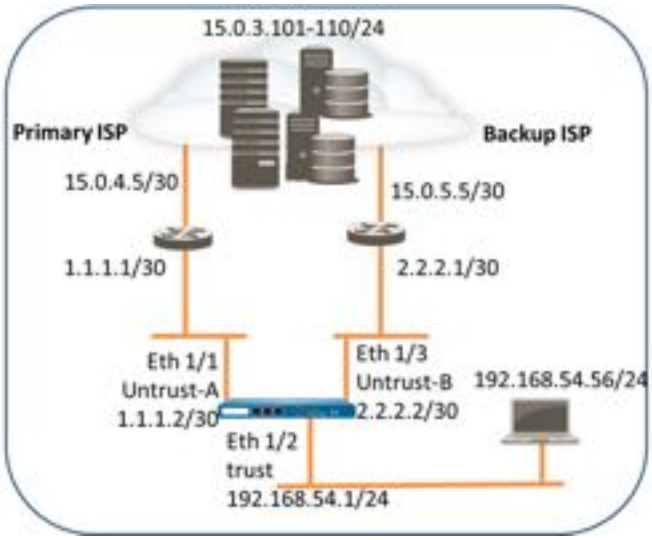


 © Palo Alto Networks, Inc.

| PBF for Outbound Access with Dual ISPs | | |
|---|---|---|
| Step 1 | Configure the ingress and the egress interfaces on the firewall.<br><br>Egress interfaces can be in the same zone. In this example we assign the egress interfaces to different zones. | 1. Select **Network > Interfaces** and then select the interface you want to configure, for example, Ethernet1/1 and Ethernet1/3.<br>The interface configuration on the firewall used in this example is as follows:<br>• Ethernet 1/1 connected to the primary ISP:<br>– Zone: ISP-East<br>– IP Address:1.1.1.2/30<br>– Virtual Router: Default<br>• Ethernet 1/3 connected to the backup ISP:<br>– Zone: ISP-West<br>– IP Address:2.2.2.2/30<br>– Virtual Router: Default<br>• Ethernet 1/2 is the ingress interface, used by the network clients to connect to the Internet:<br>– Zone: Trust<br>– IP Address:192.168. 54.1/24<br>– Virtual Router: Default<br>2. To save the interface configuration, click **OK**. |
| Step 2 | On the virtual router, add a static route to the backup ISP. | 1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog.<br>2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and specify the **Destination** IP address for which you are defining the static route. In this example, we use 0.0.0.0/0 for all traffic.<br>3. Select the **IP Address** radio button and set the **Next Hop** IP address for your router that connects to the backup Internet gateway. In this example, 2.2.2.1.<br>4. Specify a cost metric for the route. In this example, we use 10.<br> <br>5. Click **OK** twice to save the virtual router configuration. |

**PBF for Outbound Access with Dual ISPs**

Step 3 Create a PBF rule that directs traffic to the interface that is connected to the primary ISP.

Make sure to exclude traffic destined to internal servers/IP addresses from PBF. Define a negate rule so that traffic destined to internal IP addresses is not routed through the egress interface defined in the PBF rule.

1. Select **Policies > Policy Based Forwarding** and click **Add.**
2. Give the rule a descriptive **Name** in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination/Application/Service** tab, set the following:
   a. In the Destination Address section, **Add** the IP addresses or address range for servers on the internal network or create an address object for your internal servers. Select **Negate** to exclude the IP addresses or address object listed above from using this rule.
   b. In the Service section, **Add** the **service-http** and **service-https** services to allow HTTP and HTTPS traffic to use the default ports. For all other traffic that is allowed by security policy, the default route will be used.

   To forward all traffic using PBF, set the Service to **Any**.

5. In the **Forwarding** tab, specify the interface to which you want to forward traffic and enable path monitoring.
   a. To forward traffic, set the **Action** to **Forward**, and select the **Egress Interface** and specify the **Next Hop**. In this example, the egress interface is ethernet1/1, and the next hop IP address is 1.1.1.1.

© Palo Alto Networks, Inc.

**PBF for Outbound Access with Dual ISPs**

b. Enable **Monitor** and attach the default monitoring profile, to trigger a failover to the backup ISP. In this example, we do not specify a target IP address to monitor. The firewall will monitor the next hop IP address; if this IP address is unreachable the firewall will direct traffic to the default route specified on the virtual router.

c. (Required if you have asymmetric routes). Select **Enforce Symmetric Return** to ensure that return traffic from the Trust zone to the Internet is forwarded out on the same interface through which traffic ingressed from the Internet.

NAT ensures that the traffic from the Internet is returned to the correct interface/IP address on the firewall.

d. Click **OK** to save the changes.

**PBF for Outbound Access with Dual ISPs**

Step 4 Create NAT rules based on the egress interface and ISP. These rules ensure that the correct source IP address is used for outbound connections.

1. Select **Policies > NAT** and click **Add**.
2. In this example, the NAT rule we create for each ISP is as follows:

**NAT for Primary ISP**

In the **Original Packet** tab,

- **Source Zone**: Trust
- **Destination Zone**: ISP-West

In the **Translated Packet** tab, under Source Address Translation

- **Translation Type**: Dynamic IP and Port
- **Address Type**: Interface Address
- **Interface**: ethernet1/1
- **IP Address**: 1.1.1.2/30

**NAT for Backup ISP**

In the **Original Packet** tab,

- **Source Zone**: Trust
- **Destination Zone**: ISP-East

In the **Translated Packet** tab, under Source Address Translation

- **Translation Type**: Dynamic IP and Port
- **Address Type**: Interface Address
- **Interface**: ethernet1/3
- **IP Address**: 2.2.2.2/30

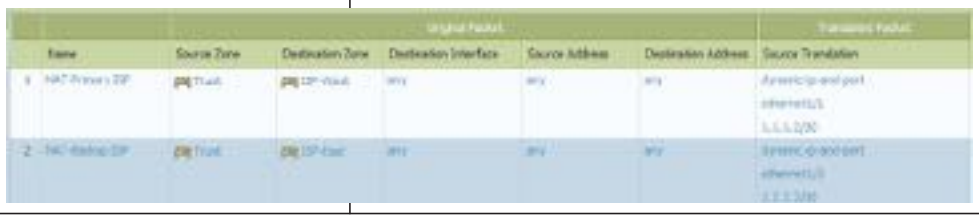

 © Palo Alto Networks, Inc.

| PBF for Outbound Access with Dual ISPs | | |
|---|---|---|
| Step 5 | Create security policy to allow outbound access to the Internet. | To safely enable applications, create a simple rule that allows access to the Internet and attach the security profiles available on the firewall.<br>1. Select **Policies > Security** and click **Add**.<br>2. Give the rule a descriptive **Name** in the **General** tab.<br>3. In the **Source** tab, set the **Source Zone** to Trust.<br>4. In the **Destination** tab, Set the **Destination Zone** to ISP-East and ISP-West.<br>5. In the **Service/ URL Category** tab, leave the default **application-default**.<br>6. In the **Actions** tab, complete these tasks:<br>a. Set the **Action Setting** to **Allow**.<br>b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection and URL filtering, under **Profile Setting.**<br>7. Under **Options**, verify that logging is enabled at the end of a session. Only traffic that matches a security rule is logged.<br> 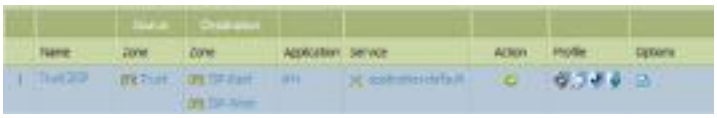 |
| Step 6 | Save the policies to the running configuration on the device. | Click **Commit**. |

**PBF for Outbound Access with Dual ISPs**

**Step 7** Verify that the PBF rule is active and that the primary ISP is used for Internet access.

1. Launch a web browser and access a web server. On the firewall check the traffic log for web-browsing activity.

2. From a client on the network, use the ping utility to verify connectivity to a web server on the Internet. and check the traffic log on the firewall.

```
C:\Users\pm-user1>ping 4.2.2.1
Pinging 4.2.2.1 with 32 bytes of data:
Reply from 4.2.2.1: bytes=32 time=34ms TTL=117
Reply from 4.2.2.1: bytes=32 time=13ms TTL=117
Reply from 4.2.2.1: bytes=32 time=25ms TTL=117
Reply from 4.2.2.1: bytes=32 time=3ms TTL=117
Ping statistics for 4.2.2.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 3ms, Maximum = 34ms, Average = 18ms
```

3. To confirm that the PBF rule is active, use the CLI command **show pbf rule all**

```
admin@PA-NGFW> show pbf rule all
Rule       ID    Rule State Action   Egress IF/VSYS  NextHop
========== === ========== ====== ============== =======
Use ISP-Pr 1     Active     Forward ethernet1/1     1.1.1.1
```

**Step 8** Verify that the failover to the backup ISP occurs and that the Source NAT is correctly applied.

1. Unplug the connection to the primary ISP.
2. Confirm that the PBF rule is inactive with the CLI command **show pbf rule all**

```
admin@PA-NGFW> show pbf rule all
Rule       ID    Rule State Action   Egress IF/VSYS  NextHop
========== === ========== ====== ============== =======
Use ISP-Pr 1   Disabled   Forward   ethernet1/1     1.1.1.1
```

3. Access a web server, and check the traffic log to verify that traffic is being forwarded through the backup ISP.

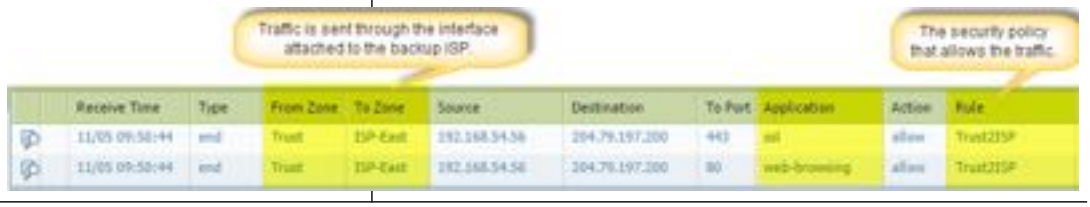


 © Palo Alto Networks, Inc.

**PBF for Outbound Access with Dual ISPs**

4. View the session details to confirm that the NAT rule is working properly.

```
admin@PA-NGFW> show session all
--------------------------------------------------------------
ID Application    State   Type Flag Src[Sport]/Zone/Proto
  (translated IP[Port]) Vsys Dst[Dport]/Zone (translated
  IP[Port])
--------------------------------------------------------------
87212 ssl ACTIVE  FLOW  NS   192.168.54.56[53236]/Trust/6
  (2.2.2.2[12896]) vsys1 204.79.197.200[443]/ISP-East
  (204.79.197.200[443])
```

5. Obtain the session identification number from the output and view the session details. Note that the PBF rule is not used and hence is not listed in the output.

```
admin@PA-NGFW> show session id 87212

Session           87212

        c2s flow:
                source:      192.168.54.56 [Trust]
                dst:         204.79.197.200
                proto:       6
                sport:       53236          dport:      443
                state:       ACTIVE         type:       FLOW
                src user:    unknown
                dst user:    unknown

        s2c flow:
                source:      204.79.197.200 [ISP-East]
                dst:         2.2.2.2
                proto:       6
                sport:       443            dport:      12896
                state:       ACTIVE         type:       FLOW
                src user:    unknown
                dst user:    unknown
start time                      : Wed Nov5 11:16:10 2014
        timeout                         : 1800 sec
        time to live                    : 1757 sec
        total byte count(c2s)           : 1918
        total byte count(s2c)           : 4333
        layer7 packet count(c2s)        : 10
        layer7 packet count(s2c)        : 7
        vsys                            : vsys1
        application                     : ssl
        rule                            : Trust2ISP
        session to be logged at end     : True
        session in session ager         : True
        session synced from HA peer     : False
        address/port translation        : source
       nat-rule                       : NAT-Backup ISP(vsys1)
        layer7 processing               : enabled
        URL filtering enabled           : True
        URL category                    : search-engines
        session via syn-cookies         : False
        session terminated on host      : False
        session traverses tunnel        : False
        captive portal session          : False
        ingress interface               : ethernet1/2
        egress interface                : ethernet1/3
        session QoS rule                : N/A (class 4)
```

# DoS Protection Against Flooding of New Sessions

The following topics describe how to configure DoS protection to better block IP addresses in order to handle high-volume attacks more efficiently.


- DoS Protection Against Flooding of New Sessions
- Configure DoS Protection Against Flooding of New Sessions
- Use the CLI to End a Single Attacking Session
- Identify Sessions That Use an Excessive Percentage of the Packet Buffer
- Discard a Session Without a Commit


 © Palo Alto Networks, Inc.

# DoS Protection Against Flooding of New Sessions

DoS protection against flooding of new sessions is beneficial against high-volume single-session and multiple-session attacks. In a single-session attack, an attacker uses a single session to target a device behind the firewall. If a Security rule allows the traffic, the session is established and the attacker initiates an attack by sending packets at a very high rate with the same source IP address and port number, destination IP address and port number, and protocol, trying to overwhelm the target. In a multiple-session attack, an attacker uses multiple sessions (or connections per second [cps]) from a single host to launch a DoS attack.

This feature defends only against DoS attacks of new sessions, that is, traffic that has not been offloaded to hardware. An offloaded attack is not protected by this feature. However, this topic describes how you can create a Security policy rule to reset the client; the attacker reinitiates the attack with numerous connections per second and is blocked by the defenses illustrated in this topic.

- Multiple-Session DoS Attack
- Single-Session DoS Attack

## Multiple-Session DoS Attack

Configure DoS Protection Against Flooding of New Sessions by configuring a DoS Protection policy rule, which determines the criteria that, when matched by incoming packets, trigger the protect action. The DoS Protection profile counts each new connection toward the Alarm Rate, Activate Rate, and Max Rate thresholds. When the incoming new connections per second exceed the Max Rate allowed, the firewall takes the action specified in the DoS Protection policy rule.

The following figure and table describe how the Security policy rules, DoS Protection policy rules and profile work together in an example.

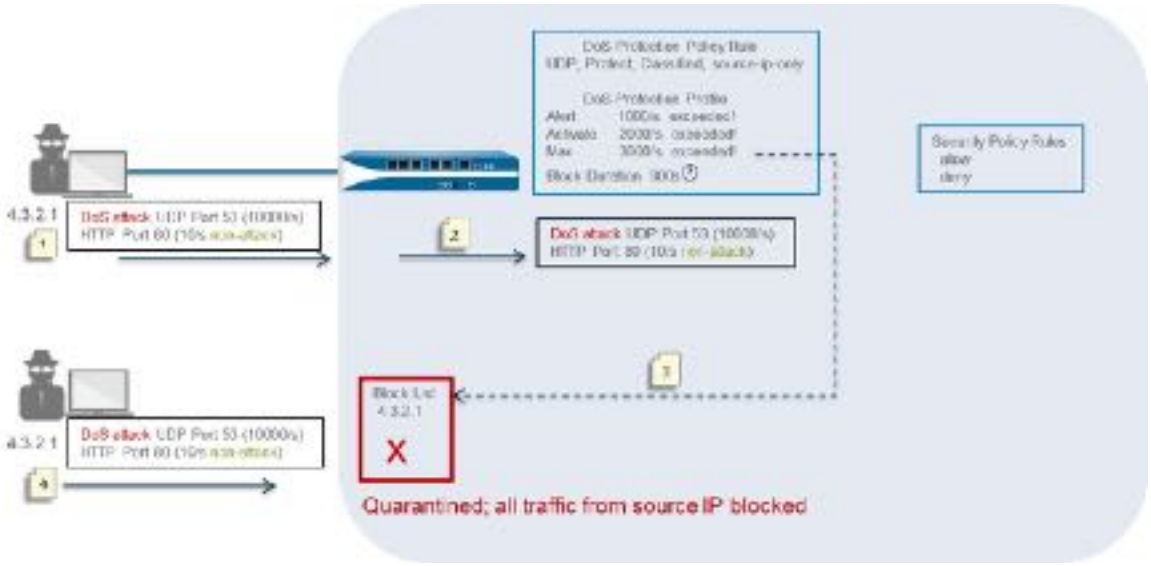

| Sequence of Events as Firewall Quarantines an IP Address | |
|---|---|
| 1 | In this example, an attacker launches a DoS attack at a rate of 10,000 new connections per second to UDP port 53. The attacker also sends 10 new connections per second to HTTP port 80. |
| 2 | The new connections match criteria in the DoS Protection policy rule, such as a source zone or interface, source IP address, destination zone or interface, destination IP address, or a service, among other settings. In this example, the policy rule specifies UDP.<br>The DoS rule also specifies the Protect action and Classified, two settings that dynamically put the DoS Protection Profile settings into effect. The DoS Protection Profile specifies that a Max Rate of 3000 packets per second is allowed. When incoming packets match the DoS rule, new connections per second are counted toward the Alert, Activate, and Max Rate thresholds.<br>You can also use a Security policy rule to block all traffic from the source IP address if you deem that address to be malicious all the time. |
| 3 | The 10,000 new connections per second exceed the Max Rate threshold. When all of the following occur:<br>• the threshold is exceeded,<br>• a Block Duration is specified, and<br>• Classified is set to includes source IP address,<br>the firewall puts the offending source IP address on the block list. |
| 4 | An IP address on the block list is in quarantine, meaning all traffic from that IP address is blocked. The firewall blocks the offending source IP address before additional attack packets reach the Security policy. |

The following figure describes in more detail what happens after an IP address that matches the DoS Protection policy rule is put on the block list. It also describes the Block Duration timer.

 © Palo Alto Networks, Inc.

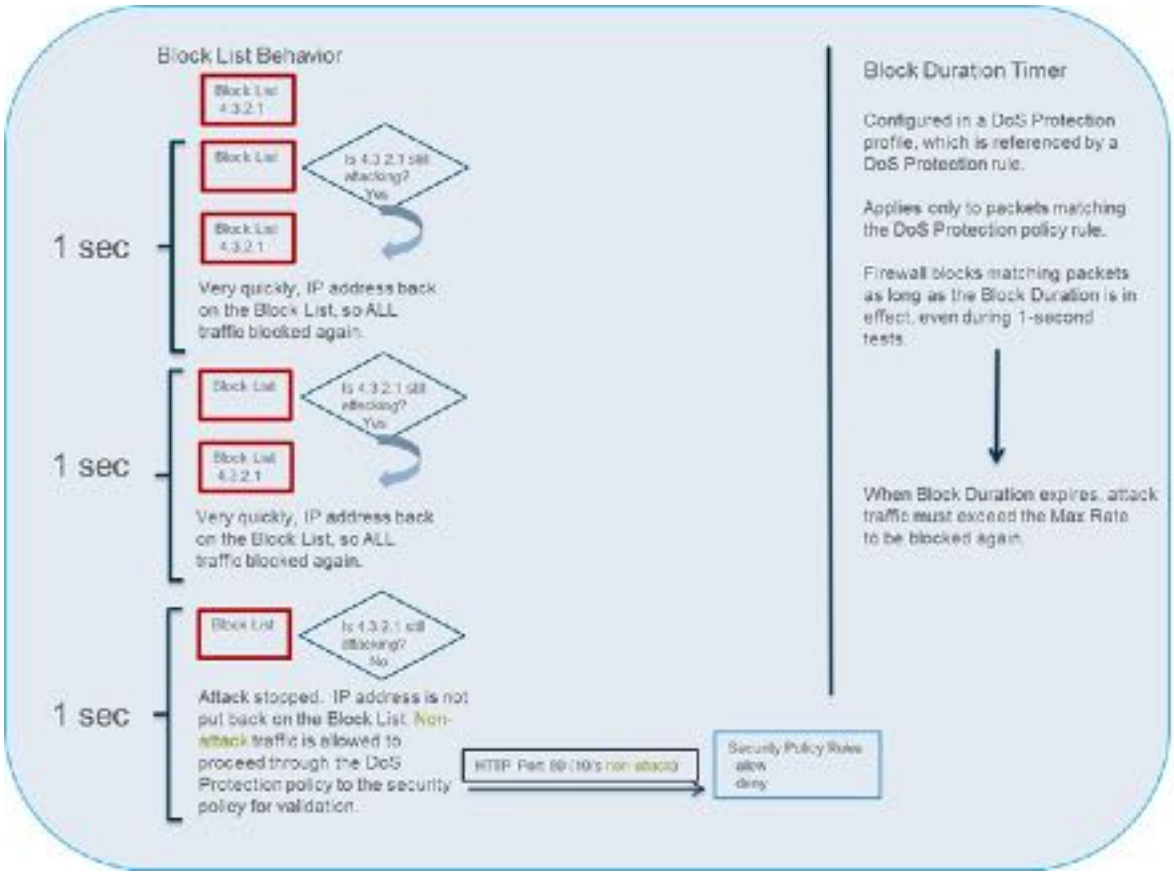


Every one second, the firewall allows the IP address to come off the Block List so that the firewall can test the traffic patterns and determine if the attack is ongoing. The firewall takes the following action:

- During this one-second test period, the firewall allows packets that do not match the DoS Protection policy criteria (HTTP traffic in this example) through the DoS Protection policy rules to the Security policy for validation. Very few packets, if any, have time to get through because the first attack packet that the firewall receives after the IP address is let off the Block List will match the DoS Protection policy criteria, quickly causing the IP address to be placed back on the block list for another second. The firewall repeats this test each second until the attack stops.
- The firewall blocks all attack traffic from going past the DoS Protection policy rules until the Block Duration expires.

When the attack stops, the firewall does not put the IP address back on the block list. The firewall allows non-attack traffic to proceed through the DoS Protection policy rules to the Security policy rules for validation. You must configure a Security policy rule because without one, an implicit deny rule denies all traffic.

The block list is based on a source zone and source address combination. This behavior allows duplicate IP addresses to exist as long as they are in different zones belonging to separate virtual routers.

The Block Duration setting in a DoS Protection profile specifies how long the firewall blocks the [offending] packets that exactly match a DoS Protection policy rule. The attack traffic remains blocked until the Block Duration expires, after which the attack traffic must again exceed the Max Rate threshold to be blocked again.

© Palo Alto Networks, Inc.

If the attacker uses multiple sessions or bots that initiate multiple attack sessions, the sessions count toward the thresholds in the DoS Protection profile without a Security policy deny rule in place. Hence, a single-session attack requires a Security policy deny rule in order for each packet to count toward the thresholds; a multiple-session attack does not.

Therefore, the DoS protection against flooding of new sessions allows the firewall to efficiently defend against a source IP address while attack traffic is ongoing and to permit non-attack traffic to pass as soon as the attack stops. Putting the offending IP address on the block list allows the DoS protection functionality to take advantage of the block list, which is designed to quarantine all activity. Quarantining the IP address from all activity protects against a modern attacker who attempts a rotating application attack, in which the attacker simply changes applications to start a new attack or uses a combination of different attacks in a hybrid DoS attack.

Beginning with PAN-OS 7.0.2, it is a change in behavior that the firewall places the attacking source IP address on the block list. When the attack stops, non-attack traffic is allowed to proceed to the Security policy rules. The attack traffic that matched the DoS Protection profile and DoS Protection policy rules remains blocked until the Block Duration expires.

## Single-Session DoS Attack

A single-session DoS attack typically will not trigger Zone or DoS Protection profiles because they are attacks that are formed after the session is created. These attacks are allowed by the Security policy because a session is allowed to be created, and after the session is created, the attack drives up the packet volume and takes down the target device.

Configure DoS Protection Against Flooding of New Sessions to protect against flooding of new sessions (single-session and multiple-session flooding). In the event of a single-session attack that is underway, additionally Use the CLI to End a Single Attacking Session.

 © Palo Alto Networks, Inc.

# Configure DoS Protection Against Flooding of New Sessions

| Configure DoS Protection Against Flooding of New Sessions | | |
|---|---|---|
| Step 1 | (Required for single-session attack mitigation or attacks that have not triggered the DoS Protection policy threshold; optional for multiple-session attack mitigation)<br><br>Configure Security policy rules to deny traffic from the attacker's IP address and allow other traffic based on your network needs. You can specify any of the match criteria in a Security policy rule, such as source IP address.<br><br> This step is one of the steps typically performed to stop an existing attack. See Use the CLI to End a Single Attacking Session. | • Set Up Basic Security Policies<br>• Components of a Security Policy Rule |

**Configure DoS Protection Against Flooding of New Sessions (Continued)**

Step 2 Configure a DoS Protection profile for flood protection.

Because flood attacks can occur over multiple protocols, as a best practice, activate protection for all of the flood types in the DoS Protection profile.

1. Select **Objects > Security Profiles > DoS Protection** and **Add** a profile **Name**.
2. Select **Classified** as the **Type**.
3. For **Flood Protection**, select the check boxes for all of the following types of flood protection:
   - **SYN Flood**
   - **UDP Flood**
   - **ICMP Flood**
   - **ICMPv6 Flood**
   - **Other IP Flood**
4. (Optional) On each of the flood tabs, change the following thresholds to suit your environment:
   - **Alarm Rate (packets/s)**—Specify the threshold rate (packets per second [pps]) above which a DoS alarm is generated. (Range is 0-2000000; default is 10000.)
   - **Activate Rate (packets/s)**—Specify the threshold rate (pps) above which a DoS response is activated. The DoS response is configured in the **Action** field of the DoS policy where this profile is referenced. When the **Activate Rate** threshold is reached, **Random Early Drop** occurs. (Range is 0-2000000; default is 10000.)
   - **Max Rate (packets/s)**—Specify the threshold rate of incoming packets per second that the firewall allows. When the threshold is exceeded, new packets that arrive are dropped and the **Action** specified in the DoS Policy rule is triggered. (Range is 2-2000000; default is 40000.)

   The default threshold values in this step are only starting points and might not be appropriate for your network. You must analyze the behavior of your network to properly set initial threshold values.
5. On each of the flood tabs, specify the **Block Duration** (in seconds), which is the length of time the firewall blocks packets that match the DoS Protection policy rule that references this profile. Specify a value greater than zero. (Range is 1-21600; default is 300.)

   Set a low Block Duration value if you are concerned that packets you incorrectly identified as attack traffic will be blocked unnecessarily.

   Set a high Block Duration value if you are more concerned about blocking volumetric attacks than you are about incorrectly blocking packets that are not part of an attack.
6. Click **OK**.


© Palo Alto Networks, Inc.

| Configure DoS Protection Against Flooding of New Sessions (Continued) | | |
|---|---|---|
| Step 3 | Configure a DoS Protection policy rule that specifies the criteria for matching the incoming traffic. | 1. Select **Policies > DoS Protection** and **Add** a **Name** on the **General** tab. The name is case-sensitive and can be a maximum of 31 characters, including letters, numbers, spaces, hyphens, and underscores.<br>2. On the **Source** tab, choose the **Type** to be a **Zone** or **Interface**, and then **Add** the zone(s) or interface(s).<br>3. (Optional) For **Source Address**, select **Any** for any incoming IP address to match the rule or **Add** an address object such as a geographical region.<br>4. (Optional) For **Source User**, select **any** or specify a user.<br>5. (Optional) Select **Negate** to match any sources except those you specify.<br>6. (Optional) On the **Destination** tab, choose the **Type** to be a **Zone** or **Interface**, and then **Add** the destination zone(s) or interface(s). For example, enter the security zone you want to protect.<br>7. (Optional) For **Destination Address**, select **Any** or enter the IP address of the device you want to protect.<br>8. (Optional) On the **Option/Protection** tab, **Add** a **Service**. Select a service or click **Service** and enter a **Name**. Select **TCP** or **UDP**. Enter a **Destination Port**. Not specifying a particular service allows the rule to match a flood of any protocol type without regard to an application-specific port.<br>9. On the **Option/Protection** tab, for **Action**, select **Protect**.<br>10. Select the **Classified** check box.<br>11. For **Profile**, select the name of the **DoS Protection** profile you created in Step 2.<br>12. For **Address**, select **source-ip-only** or **src-dest-ip-both**, which determines the type of IP address to which the rule applies. Choose the setting based on how you want the firewall to identify offending traffic.<br>• Specify **source-ip-only** if you want the firewall to classify only on the source IP address. Because attackers often test the entire network for hosts to attack, **source-ip-only** is the typical setting for a wider examination.<br>• Specify **src-dest-ip-both** if you want to protect only against DoS attacks on the server that has a specific destination address and also ensure that every source IP address will not surpass a specific connections-per-second threshold to that server.<br>13. Click **OK**. |
| Step 4 | Save the configuration. | Click **Commit**. |

## Use the CLI to End a Single Attacking Session

To mitigate a single-session DoS attack, you would still Configure DoS Protection Against Flooding of New Sessions in advance. At some point after you configure the feature, a session might be established before you realize a DoS attack (from the IP address of that session) is underway. When you see a single-session DoS attack, perform the following task to end the session, so that subsequent connection attempts from that IP address trigger the DoS protection against flooding of new sessions.

| Use the CLI to End a Single Attacking Session | |
|---|---|
| Step 1 | Identify the source IP address that is causing the attack.<br>For example, use the firewall Packet Capture feature with a destination filter to collect a sample of the traffic going to the destination IP address. Alternatively, in PAN-OS 7.0 and later, you can use ACC to filter on destination address to view the activity to the target host being attacked. |
| Step 2 | Create a DoS Protection policy rule that will block the attacker's IP address after the attack thresholds are exceeded. |
| Step 3 | Create a Security policy rule to deny the source IP address and its attack traffic. |
| Step 4 | End any existing attacks from the attacking source IP address by executing the `clear session all filter source <ip-address>` operational command.<br>Alternatively, if you know the session ID, you can execute the `clear session id <value>` command to end that session only.<br>If you use the `clear session all filter source <ip-address>` command, all sessions matching the source IP address are discarded, which can include both good and bad sessions. |

After you end the existing attack session, any subsequent attempts to form an attack session are blocked by the Security policy. The DoS Protection policy counts all connection attempts toward the thresholds. When the Max Rate threshold is exceeded, the source IP address is blocked for the Block Duration, as described in Sequence of Events as Firewall Quarantines an IP Address.

 © Palo Alto Networks, Inc.

## Identify Sessions That Use an Excessive Percentage of the Packet Buffer

When a firewall exhibits signs of resource depletion, it might be experiencing an attack that is sending an overwhelming number of packets. In such events, the firewall starts buffering inbound packets. With PAN-OS 7.0.6 and later PAN-OS 7.0 releases, you can quickly identify the sessions that are using an excessive percentage of the packet buffer and mitigate their impact by discarding them.

Perform the following task on any hardware-based firewall platform (not a VM-Series firewall) to identify, for each slot and dataplane, the packet buffer percentage used, the top five sessions using more than two percent of the packet buffer, and the source IP addresses associated with those sessions. Having that information allows you to take appropriate action.

**View Firewall Resource Usage, Top Sessions, and Session Details**

Step 1 View firewall resource usage, top sessions, and session details. Execute the following operational command in the CLI (sample output from the command follows):

```
admin@PA-7050> show running resource-monitor ingress-backlogs

-- SLOT:s1, DP:dp1 --
USAGE - ATOMIC: 92%  TOTAL: 93%

TOP SESSIONS:
SESS-ID      PCT   GRP-ID   COUNT
6            92%   1        156
                   7        1732

SESSION DETAILS
SESS-ID PROTO SZONE SRC           SPORT  DST       DPORT  IGR-IF    EGR-IF       APP
6       6     trust 192.168.2.35 55653  10.1.8.89 80  ethernet1/21 ethernet1/22
undecided
```

The command displays a maximum of the top five sessions that each use 2% or more of the packet buffer.

The sample output above indicates that Session 6 is using 92% of the packet buffers with TCP packets (protocol 6) coming from source IP address 192.168.2.35.

- **GRP-ID**—indicates an internal stage of processing packets.
- **COUNT**—indicates how many packets are in that GRP-ID for that session.
- **SESS-ID**—indicates the global session ID that is used in all other `show session` commands. The global session ID is unique within the firewall.
- **APP**—indicates the App-ID extracted from the Session information, which can help you determine whether the traffic is legitimate. For example, if packets use a common TCP or UDP port but the CLI output indicates an APP of `undecided`, the packets are possibly attack traffic. The APP is `undecided` when Application IP Decoders cannot get enough information to determine the application. An APP of `unknown` indicates that Application IP Decoders cannot determine the application; a session of `unknown` APP that uses a high percentage of the packet buffer is also suspicious.

To restrict the display output:

On a PA-7000 Series platform, you can limit output to a slot, a dataplane, or both. For example:

```
admin@PA-7050> show running resource-monitor ingress-backlogs slot s1
admin@PA-7050> show running resource-monitor ingress-backlogs slot s1 dp dp1
```

On a PA-5000 Series platform, you can limit output to a dataplane. For example:

```
admin@PA-5060> show running resource-monitor ingress-backlogs dp dp1
```

**View Firewall Resource Usage, Top Sessions, and Session Details (Continued)**

Step 2 Use the command output to determine whether the source at the source IP address using a high percentage of the packet buffer is sending legitimate or attack traffic.

In the sample output above, a single-session attack is likely occurring. A single session (Session ID 6) is using 92% of the packet buffer for Slot 1, DP 1, and the application at that point is `undecided`.

- If you determine a single user is sending an attack and the traffic is in the slow path, you can Use the CLI to End a Single Attacking Session. At a minimum, you can Configure DoS Protection Against Flooding of New Sessions.
- If the traffic is offloaded to hardware (the traffic is in the fast path), clearing the session does not help because the software then has to handle the barrage of packets. You should Discard a Session Without a Commit instead.

To see whether a session is offloaded or not, use the `show session id <session-id>` operational command in the CLI as shown in the following example. The `layer7 processing value` indicates `completed` for sessions offloaded or `enabled` for sessions not offloaded.

```
admin@PA-5060> show session id 68088184

Session       68088184

        c2s flow:
                source:      1.1.42.15 [trust]
                dst:         1.2.27.99
                proto:       6
                sport:       55993           dport:      6881
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
                offload:     Yes

        s2c flow:
                source:      1.2.27.99 [untrust]
                dst:         1.1.42.15
                proto:       6
                sport:       6881            dport:      55993
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
                offload:     Yes

        DP                                   : 2
        index(local):                        : 979328
        start time                           : Tue Oct 27 14:20:09 2015
        timeout                              : 1200 sec
        time to live                         : 1167 sec
        total byte count(c2s)                : 270
        total byte count(s2c)                : 270
        layer7 packet count(c2s)             : 3
        layer7 packet count(s2c)             : 3
        vsys                                 : vsys1
        application                          : bittorrent
        rule                                 : rule1
        session to be logged at end          : True
        session in session ager              : True
        session updated by HA peer           : False
        layer7 processing                    : completed
        URL filtering enabled                : False
        session via syn-cookies              : False
        session terminated on host           : False
        session traverses tunnel             : False
        captive portal session               : False
        ingress interface                    : ethernet1/21
        egress interface                     : ethernet1/22
        session QoS rule                     : N/A (class 4)
        tracker stage l7proc                 : ctd decoder bypass
        end-reason                           : unknown
```

 © Palo Alto Networks, Inc.

## Discard a Session Without a Commit

With PAN-OS 7.0.6 and later PAN-OS 7.0 releases, you can perform this task to permanently discard a session, such as a session that is overloading the packet buffer. No commit is required; the session is discarded immediately after executing the command. The commands you use for this task apply to both offloaded and non-offloaded sessions.

| **Discard a Session Without a Commit** | |
|---|---|
| Step 1 | In the CLI, execute the following operational command on any hardware platform:<br>`admin@PA-7050>` **request session-discard [timeout <seconds>] [reason <reason-string>] id <session-id>**<br>The default timeout is 3600 seconds. |
| Step 2 | Verify that sessions have been discarded:<br>`admin@PA-7050>` **show session all filter state discard** |

 © Palo Alto Networks, Inc.

# Virtual Systems

This topic describes virtual systems, their benefits, typical use cases, and how to configure them. It also provides links to other topics where virtual systems are documented as they function with other features.


- Virtual Systems Overview
- Communication Between Virtual Systems
- Shared Gateway
- Configure Virtual Systems
- Configure Inter-Virtual System Communication within the Firewall
- Configure a Shared Gateway
- Service Routes for Virtual Systems
- Customize Service Routes for a Virtual System
- DNS Resolution—Three Use Cases
- Virtual System Functionality with Other Features


© Palo Alto Networks, Inc.

# Virtual Systems Overview

Virtual systems are separate, logical firewall instances within a single physical Palo Alto Networks firewall. Rather than using multiple firewalls, managed service providers and enterprises can use a single pair of firewalls (for high availability) and enable virtual systems on them. Each virtual system (vsys) is an independent, separately-managed firewall with its traffic kept separate from the traffic of other virtual systems.

This topic includes the following:


- Virtual System Components and Segmentation
- Benefits of Virtual Systems
- Use Cases for Virtual Systems
- Platform Support and Licensing for Virtual Systems
- Administrative Roles for Virtual Systems
- Shared Objects for Virtual Systems


## Virtual System Components and Segmentation

A virtual system is an object that creates an administrative boundary, as shown in the following figure.

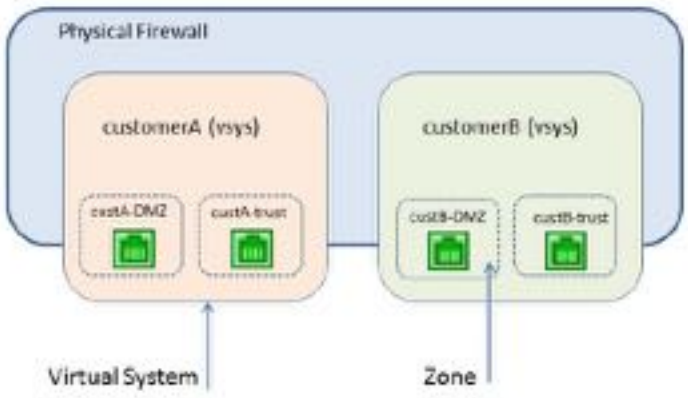


A virtual system consists of a set of physical and logical interfaces and subinterfaces (including VLANs and virtual wires), virtual routers, and security zones. You choose the deployment mode(s) (any combination of virtual wire, Layer 2, or Layer 3) of each virtual system. By using virtual systems, you can segment any of the following:

- Administrative access
- The management of all policies (security, NAT, QoS, policy-based forwarding, decryption, application override, captive portal, and DoS protection)
- All objects (such as address objects, application groups and filters, dynamic block lists, security profiles, decryption profiles, custom objects, etc.)
- User-ID

 © Palo Alto Networks, Inc.

- Certificate management
- Server profiles
- Logging, reporting, and visibility functions

Virtual systems affect the security functions of the firewall, but virtual systems alone do not affect networking functions such as static and dynamic routing. You can segment routing for each virtual system by creating one or more virtual routers for each virtual system, as in the following use cases:

- If you have virtual systems for departments of one organization, and the network traffic for all of the departments is within a common network, you can create a single virtual router for multiple virtual systems.
- If you want routing segmentation and each virtual system's traffic must be isolated from other virtual systems, you can create one or more virtual routers for each virtual system.

## Benefits of Virtual Systems

Virtual systems provide the same basic functions as a physical firewall, along with additional benefits:

- **Segmented administration**—Different organizations (or customers or business units) can control (and monitor) a separate firewall instance, so that they have control over their own traffic without interfering with the traffic or policies of another firewall instance on the same physical device.
- **Scalability**—After the physical firewall is configured, adding or removing customers or business units can be done efficiently. An ISP, managed security service provider, or enterprise can provide different security services to each customer.
- **Reduced capital and operational expenses**—Virtual systems eliminate the need to have multiple physical firewalls at one location because virtual systems co-exist on one firewall. By not having to purchase multiple firewalls, an organization can save on the hardware expense, electric bills, and rack space, and can reduce maintenance and management expenses.

## Use Cases for Virtual Systems

There are many ways to use virtual systems in a network. One common use case is for an ISP or a managed security service provider (MSSP) to deliver services to multiple customers with a single firewall. Customers can choose from a wide array of services that can be enabled or disabled easily. The firewall's role-based administration allows the ISP or MSSP to control each customer's access to functionality (such as logging and reporting) while hiding or offering read-only capabilities for other functions.

Another common use case is within a large enterprise that requires different firewall instances because of different technical or confidentiality requirements among multiple departments. Like the above case, different groups can have different levels of access while IT manages the firewall itself. Services can be tracked and/or billed back to departments to thereby make separate financial accountability possible within an organization.

© Palo Alto Networks, Inc.

## Platform Support and Licensing for Virtual Systems

Virtual systems are supported on the PA-2000, PA-3000, PA-4000, PA-5000, and PA-7000 Series firewalls. Each firewall series supports a base number of virtual systems; the number varies by platform. A Virtual Systems license is required in the following cases:

- To support multiple virtual systems on PA-2000 or PA-3000 Series firewalls.
- To create more than the base number of virtual systems supported on a platform.

For license information, see Activate Licenses and Subscriptions. For the base and maximum number of virtual systems supported, see Compare Firewalls tool.

Multiple virtual systems are not supported on the PA-200, PA-500 or VM-Series firewalls.

## Administrative Roles for Virtual Systems

A **superuser** administrator can create virtual systems and add a **Device Administrator**, **vsysadmin**, or **vsysreader**. A **Device Administrator** can access all virtual systems, but cannot add administrators. The two types of virtual system administrative roles are:

- **vsysadmin**—Grants full access to a virtual system.
- **vsysreader**—Grants read-only access to a virtual system.

A virtual system administrator can view logs of only the virtual systems assigned to that administrator. Someone with **superuser** or **Device Admin** permission can view all of the logs or select a virtual system to view.

Persons with **vsysadmin** permission can commit configurations for only the virtual systems assigned to them.

## Shared Objects for Virtual Systems

If your administrator account extends to multiple virtual systems, you can choose to configure objects (such as an address object) and policies for a specific virtual system or as shared objects, which apply to all of the virtual systems on the firewall. If you try to create a shared object with the same name and type as an existing object in a virtual system, the virtual system object is used.

 © Palo Alto Networks, Inc.

# Communication Between Virtual Systems

There are two typical scenarios where communication between virtual systems (inter-vsys traffic) is desirable. In a multi-tenancy environment, communication between virtual systems can occur by having traffic leave the firewall, go through the Internet, and re-enter the firewall. In a single organization environment, communication between virtual systems can remain within the firewall. This section discusses both scenarios.


- Inter-VSYS Traffic That Must Leave the Firewall
- Inter-VSYS Traffic That Remains Within the Firewall
- Inter-VSYS Communication Uses Two Sessions

## Inter-VSYS Traffic That Must Leave the Firewall

An ISP that has multiple customers on a firewall (known as multi-tenancy) can use a virtual system for each customer, and thereby give each customer control over its virtual system configuration. The ISP grants **vsysadmin** permission to customers. Each customer's traffic and management are isolated from the others. Each virtual system must be configured with its own IP address and one or more virtual routers in order to manage traffic and its own connection to the Internet.

If the virtual systems need to communicate with each other, that traffic goes out the firewall to another Layer 3 routing device and back to the firewall, even though the virtual systems exist on the same physical firewall, as shown in the following figure.

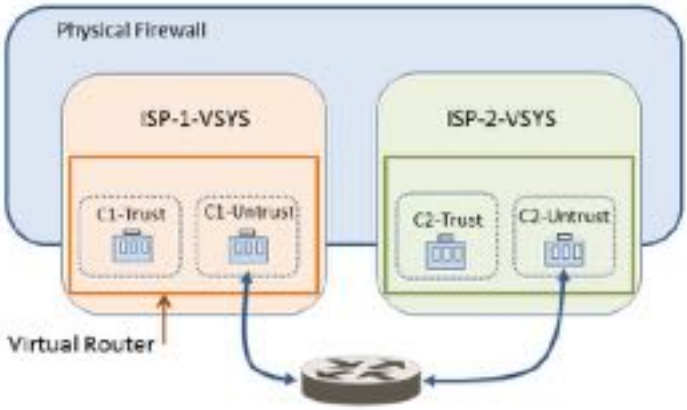


 © Palo Alto Networks, Inc.

# Inter-VSYS Traffic That Remains Within the Firewall

Unlike the preceding multi-tenancy scenario, virtual systems on a firewall can be under the control of a single organization. The organization wants to both isolate traffic between virtual systems and allow communications between virtual systems. This common use case arises when the organization wants to provide departmental separation and still have the departments be able to communicate with each other or connect to the same network(s). In this scenario, the inter-vsys traffic remains within the firewall, as described in the following topics:


- External Zone
- External Zones and Security Policies For Traffic Within a Firewall


## External Zone

The communication desired in the use case above is achieved by configuring security policies that point to or from an *external* zone. An external zone is a security object that is associated with a specific virtual system that it can reach; the zone is external to the virtual system. A virtual system can have only one external zone, regardless of how many security zones the virtual system has within it. External zones are required to allow traffic between zones in different virtual systems, without the traffic leaving the firewall.

The virtual system administrator configures the security policies needed to allow traffic between two virtual systems. Unlike security zones, an external zone is not associated with an interface; it is associated with a virtual system. The security policy allows or denies traffic between the security (internal) zone and the external zone.

Because external zones do not have interfaces or IP addresses associated with them, some zone protection profiles are not supported on external zones.

Remember that each virtual system is a separate instance of a firewall, which means that each packet moving between virtual systems is inspected for security policy and App-ID evaluation.

## External Zones and Security Policies For Traffic Within a Firewall

In the following example, an enterprise has two separate administrative groups: the departmentA and departmentB virtual systems. The following figure shows the external zone associated with each virtual system, and traffic flowing from one trust zone, out an external zone, into an external zone of another virtual system, and into its trust zone.

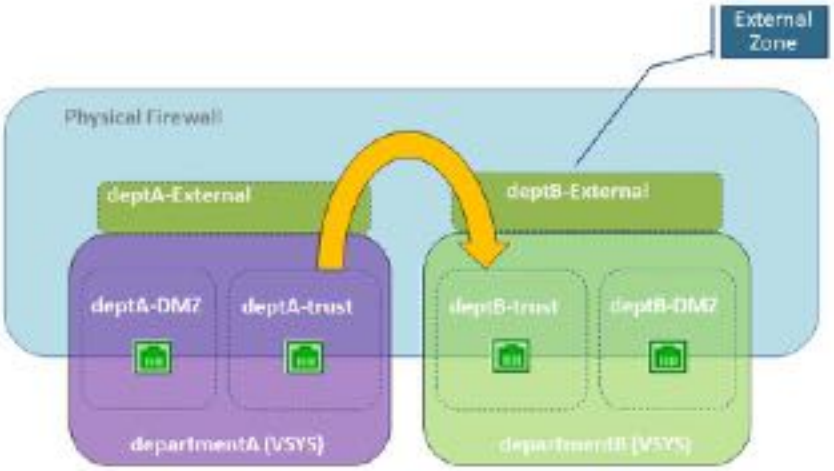


In order to create external zones, the device administrator must configure the virtual systems so that they are *visible* to each other. External zones do not have security policies between them because their virtual systems are visible to each other.

To communicate between virtual systems, the ingress and egress interfaces on the firewall are either assigned to a single virtual router or else they are connected using inter-virtual router static routes. The simpler of these two approaches is to assign all virtual systems that must communicate with each other to a single virtual router.

There might be a reason that the virtual systems need to have their own virtual router, for example, if the virtual systems use overlapping IP address ranges. Traffic can be routed between the virtual systems, but each virtual router must have static routes that point to the other virtual router(s) as the next hop.

Referring to the scenario in the figure above, we have an enterprise with two administrative groups: departmentA and departmentB. The departmentA group manages the local network and the DMZ resources. The departmentB group manages traffic in and out of the sales segment of the network. All traffic is on a local network, so a single virtual router is used. There are two external zones configured for communication between the two virtual systems. The departmentA virtual system has three zones used in security policies: deptA-DMZ, deptA-trust, and deptA-External. The departmentB virtual system also has three zones: deptB-DMZ, deptB-trust, and deptB-External. Both groups can control the traffic passing through their virtual systems.

In order to allow traffic from deptA-trust to deptB-trust, two security policies are required. In the following figure, the two vertical arrows indicate where the security policies (described below the figure) are controlling traffic.

 © Palo Alto Networks, Inc.

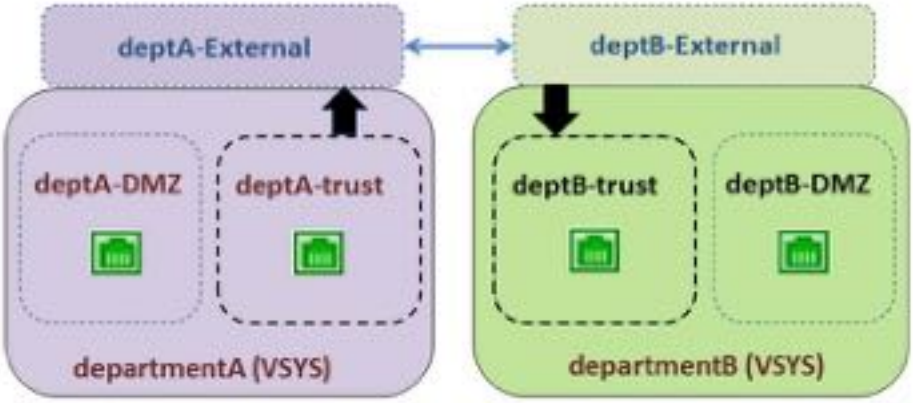


- Security Policy 1: In the preceding figure, traffic is destined for the deptB-trust zone. Traffic leaves the deptA-trust zone and goes to the deptA-External zone. A security policy must allow traffic from the source zone (deptA-trust) to the destination zone (deptA-External). A virtual system allows any policy type to be used for this traffic, including NAT.

  No policy is needed between external zones because traffic sent to an external zone appears in and has automatic access to the other external zones that are visible to the original external zone.

- Security Policy 2: In the preceding figure, the traffic from deptB-External is still destined to the deptB-trust zone, and a security policy must be configured to allow it. The policy must allow traffic from the source zone (deptB-External) to the destination zone (deptB-trust).

The departmentB virtual system could be configured to block traffic from the departmentA virtual system, and vice versa. Like traffic from any other zone, traffic from external zones must be explicitly allowed by policy to reach other zones in a virtual system.

In addition to external zones being required for inter-virtual system traffic that does not leave the firewall, external zones are also required if you configure a Shared Gateway, in which case the traffic is intended to leave the firewall.

## Inter-VSYS Communication Uses Two Sessions

It is helpful to understand that communication between two virtual systems uses two sessions, unlike the one session used for a single virtual system. Let's compare the scenarios.

Scenario 1—Vsys1 has two zones: trust1 and untrust1. A host in the trust1 zone initiates traffic when it needs to communicate with a device in the untrust1 zone. The host sends traffic to the firewall, and the firewall creates a new session for source zone trust1 to destination zone untrust1. Only one session is needed for this traffic.

Scenario 2—A host from vsys1 needs to access a server on vsys2. A host in the trust1 zone initiates traffic to the firewall, and the firewall creates the first session: source zone trust1 to destination zone untrust1. Traffic is routed to vsys2, either internally or externally. Then the firewall creates a second session: source zone untrust2 to destination zone trust2. Two sessions are needed for this inter-vsys traffic.

 © Palo Alto Networks, Inc.

# Shared Gateway

This topic includes the following information about shared gateways:

- External Zones and Shared Gateway
- Networking Considerations for a Shared Gateway

## External Zones and Shared Gateway

A shared gateway is an interface that multiple virtual systems share in order to communicate over the Internet. Each virtual system requires an External Zone, which acts as an intermediary, for configuring security policies that allow or deny traffic from the virtual system's internal zone to the shared gateway.

The shared gateway uses a single virtual router to route traffic for all virtual systems. A shared gateway is used in cases when an interface does not need a full administrative boundary around it, or when multiple virtual systems must share a single Internet connection. This second case arises if an ISP provides an organization with only one IP address (interface), but multiple virtual systems need external communication.

Unlike the behavior between virtual systems, security policy and App-ID evaluations are not performed between a virtual system and a shared gateway. That is why using a shared gateway to access the Internet involves less overhead than creating another virtual system to do so.

In the following figure, three customers share a firewall, but there is only one interface accessible to the Internet. Creating another virtual system would add the overhead of App-ID and security policy evaluation for traffic being sent to the interface through the added virtual system. To avoid adding another virtual system, the solution is to configure a shared gateway, as shown in the following diagram.

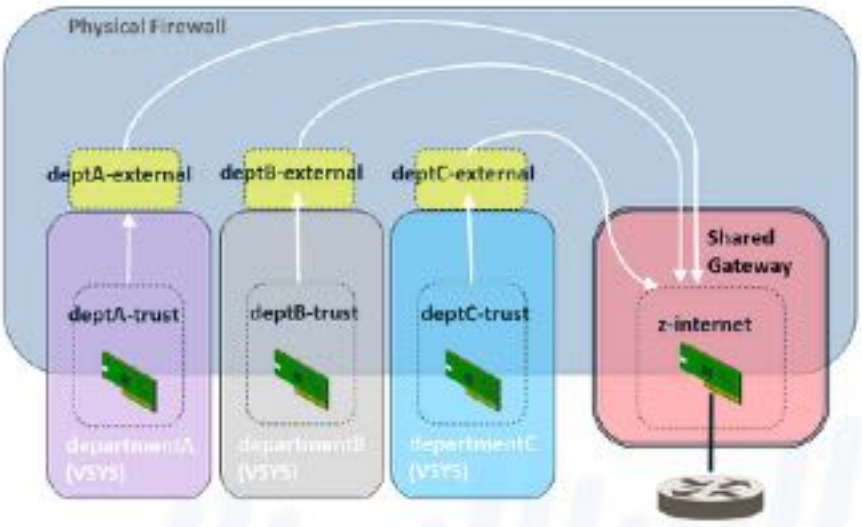


The shared gateway has one globally-routable IP address used to communicate with the outside world. Interfaces in the virtual systems have IP addresses too, but they can be private, non-routable IP addresses.

You will recall that an administrator must specify whether a virtual system is visible to other virtual systems. Unlike a virtual system, a shared gateway is always visible to all of the virtual systems on the firewall.

A shared gateway ID number appears as **sg<ID>** on the web interface. It is recommended that you name your shared gateway with a name that includes its ID number.

When you add objects such as zones or interfaces to a shared gateway, the shared gateway appears as an available virtual system in the vsys drop-down menu.

A shared gateway is a limited version of a virtual system; it supports NAT and policy-based forwarding (PBF), but does not support security, DoS policies, QoS, decryption, application override, or captive portal policies.

## Networking Considerations for a Shared Gateway

Keep the following in mind while you are configuring a shared gateway.

- The virtual systems in a shared gateway scenario access the Internet through the shared gateway's physical interface, using a single IP address. If the IP addresses of the virtual systems are not globally routable, configure source NAT to translate those addresses to globally-routable IP addresses.
- A virtual router routes the traffic for all of the virtual systems through the shared gateway.
- The default route for the virtual systems should point to the shared gateway.
- Security policies must be configured for each virtual system to allow the traffic between the internal zone and external zone, which is visible to the shared gateway.
- A device administrator should control the virtual router, so that no member of a virtual system can affect the traffic of other virtual systems.
- Within a Palo Alto Networks firewall, a packet may hop from one virtual system to another virtual system or a shared gateway. A packet may not traverse more than two virtual systems or shared gateways. For example, a packet cannot go from one virtual system to a shared gateway to a second virtual system within the firewall.

To save configuration time and effort, consider the following advantages of a shared gateway:

- Rather than configure NAT for multiple virtual systems associated with a shared gateway, you can configure NAT for the shared gateway.
- Rather than configure policy-based routing (PBR) for multiple virtual systems associated with a shared gateway, you can configure PBR for the shared gateway.

 © Palo Alto Networks, Inc.

# Service Routes for Virtual Systems

The firewall uses the MGT interface (by default) to access external services, such as DNS servers, software updates, and software licenses. An alternative to using the MGT interface is to configure a data port (a regular interface) to access these services. The path from the interface to the service on a server is known as a *service route*. Service routes can be configured for the firewall or for individual virtual systems. Each service allows redirection of management services to the respective virtual system owner through one of the interfaces associated with that virtual system.

The ability to configure service routes per virtual system provides the flexibility to customize service routes for numerous tenants or departments on a single firewall. The service packets exit the firewall on a port that is assigned to a specific virtual system, and the server sends its response to the configured source interface and source IP address. Any virtual system that does not have a service route configured for a particular service inherits the interface and IP address that are set globally for that service.

- Use Cases for Service Routes for a Virtual System
- PA-7000 Series Firewall LPC Support for Per-Virtual System Paths to Logging Servers
- DNS Proxy Object
- DNS Server Profile
- Multi-Tenant DNS Deployments

To configure service routes for a virtual system, see Customize Service Routes for a Virtual System.

## Use Cases for Service Routes for a Virtual System

One use case for configuring service routes at the virtual system level is when a large customer (such as an ISP) needs to support multiple individual tenants on a single Palo Alto Networks firewall. The ISP has configured virtual systems on the firewall, and wants to have separate service routes for each virtual system, rather than services routes configured at the global level. Each tenant requires service route capabilities so that it can customize service route parameters for DNS, email, Kerberos, LDAP, NetFlow, RADIUS, SNMP trap, syslog, TACACS+, User-ID Agent, and VM Monitor.

Another use case is an IT organization that wants to provide full autonomy to groups that set servers for services. Each group can have a virtual system and define its own service routes.

If **Multi Virtual System Capability** is enabled, any virtual system that does not have specific service routes configured inherits the global service and service route settings for the device.

An organization can have multiple virtual systems, but use a global service route for a service rather than different service routes for each virtual system. For example, the firewall can use a shared email server to originate email alerts to its virtual systems.

A firewall with multiple virtual systems must have interfaces and subinterfaces with non-overlapping IP addresses.

A per-virtual system service route for SNMP traps or for Kerberos is for IPv4 only.

You can select a virtual router for a service route in a virtual system; you cannot select the egress interface. After you select the virtual router and the firewall sends the packet from the virtual router, the firewall selects the egress interface based on the destination IP address. Therefore:

- If a virtual system has multiple virtual routers, packets to all of the servers for a service must egress out of only one virtual router.
- A packet with an interface source address may egress a different interface, but the return traffic would be on the interface that has the source IP address, creating asymmetric traffic.

 © Palo Alto Networks, Inc.

## PA-7000 Series Firewall LPC Support for Per-Virtual System Paths to Logging Servers

For Traffic, HIP Match, Threat, and Wildfire log types, the PA-7000 Series firewall does not use service routes for SNMP Trap, syslog and email services. Instead, the PA-7000 Series firewall Log Processing Card (LPC) supports virtual system-specific paths from LPC subinterfaces to an on-premise switch to the respective service on a server. For System and Config logs, the PA-7000 Series firewall uses global service routes, and not the LPC.

In other Palo Alto Networks platforms, the dataplane sends logging service route traffic to the management plane, which sends the traffic to logging servers. In the PA-7000 Series firewall, each LPC has only one interface, and data planes for multiple virtual systems send logging server traffic (types mentioned above) to the PA-7000 Series firewall LPC. The LPC is configured with multiple subinterfaces, over which the platform sends the logging service traffic out to a customer's switch, which can be connected to multiple logging servers.

Each LPC subinterface can be configured with a subinterface name and a dotted subinterface number. The subinterface is assigned to a virtual system, which is configured for logging services. The other service routes on a PA-7000 Series firewall function similarly to service routes on other Palo Alto Networks platforms.

To configure the LPC for per-virtual system logging services, see Configure a PA-7000 Series Firewall for Logging Per Virtual System. For information about the LPC itself, see the PA-7000 Series Hardware Reference Guide.

## DNS Proxy Object

Domain Name System (DNS) servers perform the service of resolving a domain name to an IP address, and vice versa. DNS proxy is a role in which the firewall is an intermediary between DNS clients and servers; it acts as a DNS server itself by resolving queries from its DNS proxy cache. If the domain name is not found in the DNS proxy cache, the firewall searches for a match to the domain name among the entries in the specific DNS proxy object (on the interface on which the DNS query arrived), and forwards the query to a DNS server based on the match results. If no match is found, the default DNS servers are used.

A DNS proxy object is where you configure the settings that determine how the firewall functions as a DNS proxy. You can assign a DNS proxy object to a single virtual system or it can be shared among all virtual systems.

- If the DNS proxy object is for a virtual system, you can specify a DNS Server Profile, which specifies the primary and secondary DNS server addresses, along with other information. The DNS server profile simplifies configuration.
- If the DNS proxy object is shared, you must specify at least the primary address of a DNS server.

When configuring tenants with DNS services, each tenant should have its own DNS proxy defined, which keeps the tenant's DNS service separate from other tenants' services.

In the proxy object, you specify the interfaces for which the firewall is acting as DNS proxy. The DNS proxy for the interface does not use the service route; responses to the DNS requests are always sent to the interface assigned to the virtual router where the DNS request arrived.

You can supply the DNS proxy with static FQDN-to-address mappings. You can create DNS proxy rules that control to which DNS server the specified domain name queries are directed. A DNS proxy has other options; to configure a DNS proxy, see Configure a DNS Proxy Object. A maximum of 256 DNS proxy objects can be configured on a firewall.

 © Palo Alto Networks, Inc.

## DNS Server Profile

To simplify configuration for a virtual system, a DNS server profile allows you to specify the virtual system that is being configured, an inheritance source or the primary and secondary IP addresses for DNS servers, and a source interface and source address (service route) that will be used in packets sent to the DNS server. The source interface determines the virtual router, which has a route table. The destination IP address is looked up in the routing table of the virtual router where the source interface is assigned. It is possible that the result of the destination IP egress interface differs from the source interface. The packet would egress out of the destination IP egress interface determined by the route table lookup, but the source IP address would be the address configured. The source address is used as the destination address in the reply from the DNS server.

The virtual system report and virtual system server profile send their queries to the DNS server specified for the virtual system, if there is one. (The DNS server used is defined in **Device > Virtual Systems > General > DNS Proxy**.) If there is no DNS server specified for the virtual system, the DNS server specified for the device is queried.

A DNS server profile is for a virtual system only; it is not for a global **Shared** location. To configure a DNS server profile, see Configure a DNS Server Profile.

For more information on DNS server profiles, see DNS Resolution—Three Use Cases.

## Multi-Tenant DNS Deployments

There are three use cases for multi-tenant DNS deployments:

- **Global Management DNS Resolution**—The firewall needs DNS resolution for its own purposes, for example, when the request is coming from the management plane to resolve an FQDN in a security policy. The firewall uses the service route to get to a DNS server because there is no incoming virtual router. The DNS server is configured in **Device > Setup > Services > Global**, and **Servers** are configured by entering a primary and secondary DNS server.
- **Policy and Report FQDN Resolution for a Virtual System**—For DNS queries that need to be resolved from a security policy or a report, you can specify a set of DNS servers specific to the virtual system (tenant) or you can default to the global DNS servers. If your use case requires a different set of DNS servers per virtual system, the DNS server is configured in **Device > Virtual Systems > General > DNS Proxy**. The DNS proxy object is configured in **Network > DNS Proxy**. The resolution is specific to the virtual system to which the DNS proxy is assigned. If you don't have specific DNS servers applicable to this virtual system and want to use the global DNS setting, the global DNS servers take precedence.
- **Dataplane DNS Resolution for a Virtual System**—This method is also known as a Network Request for DNS Resolution. The tenant's virtual system can be configured so that specified domain names are resolved on the tenant's DNS server in its network. This method supports *split DNS*, meaning that the tenant can also use its own ISP DNS servers for the remaining DNS queries not resolved on its own server. DNS Proxy rules control the split DNS; the tenant's domain redirects DNS requests to its DNS servers, which are configured in a DNS server profile. The DNS server profile has primary and secondary DNS servers designated, and also DNS service routes for IPv4 and IPv6, which override the default DNS settings.

For more information on DNS deployments, see DNS Resolution—Three Use Cases.

 © Palo Alto Networks, Inc.

# Configure Virtual Systems

Creating a virtual system requires that you have the following:

- A **superuser** administrative role.
- An interface configured.
- A Virtual Systems license if you are configuring a PA-2000 or PA-3000 Series firewall, or if you are creating more than the base number of virtual systems supported on the platform. See Platform Support and Licensing for Virtual Systems.

| **Configure a Virtual System** | |
|---|---|
| Step 1 Enable virtual systems. | 1. Select **Device > Setup > Management** and edit the **General Settings**.<br>2. Select the **Multi Virtual System Capability** check box and click **OK**. This action triggers a commit if you approve it.<br>Only after enabling virtual systems will the **Device** tab display the **Virtual Systems** and **Shared Gateways** options. |
| Step 2 Create a virtual system. | 1. Select **Device > Virtual Systems**, click **Add** and enter a virtual system **ID**, which is appended to "vsys" (range is 1-255).<br> The default ID is 1, which makes the default virtual system **vsys1**. This default appears even on platforms that do not support multiple virtual systems.<br>2. Check the **Allow forwarding of decrypted content** check box if you want to allow the firewall to forward decrypted content to an outside service. For example, you must enable this option for the firewall to be able to send decrypted content to WildFire for analysis.<br>3. Enter a descriptive **Name** for the virtual system. A maximum of 31 alphanumeric, space, and underscore characters is allowed. |

| Configure a Virtual System | | |
|---|---|---|
| Step 3 | Assign interfaces to the virtual system.<br>The virtual routers, vwires, or VLANs can either be configured already or you can configure them later, at which point you specify the virtual system associated with each. The procedure to configure a virtual router, for example, is in Step 6 below. | 1. On the **General** tab, select a **DNS Proxy** object if you want to apply DNS proxy rules to the interface.<br>2. In the **Interfaces** field, click **Add** to enter the interfaces or subinterfaces to assign to the virtual system. An interface can belong to only one virtual system.<br>3. Do any of the following, based on the deployment type(s) you need in the virtual system:<br>• In the **VLANs** field, click **Add** to enter the VLAN(s) to assign to the vsys.<br>• In the **Virtual Wires** field, click **Add** to enter the virtual wire(s) to assign to the vsys.<br>• In the **Virtual Routers** field, click **Add** to enter the virtual router(s) to assign to the vsys.<br>4. In the **Visible Virtual System** field, check all virtual systems that should be made visible to the virtual system being configured. This is required for virtual systems that need to communicate with each other.<br>In a multi-tenancy scenario where strict administrative boundaries are required, no virtual systems would be checked.<br>5. Click **OK**. |
| Step 4 | (Optional) Limit the resource allocations for sessions, rules, and VPN tunnels allowed for the virtual system. The flexibility of being able to allocate limits per virtual system allows you to effectively control device resources. | 1. On the **Resource** tab, optionally set limits for a virtual system. There are no default values.<br>• **Sessions Limit**—Range is 1-262144.<br>• **Security Rules**—Range is 0-2500.<br>• **NAT Rules**—Range is 0-3000.<br>• **Decryption Rules**—Range is 0-250.<br>• **QoS Rules**—Range is 0-1000.<br>• **Application Override Rules**—Range is 0-250.<br>• **Policy Based Forwarding Rules**—Range is 0-500.<br>• **Captive Portal Rules**—Range is 0-1000.<br>• **DoS Protection Rules**—Range is 0-1000.<br>• **Site to Site VPN Tunnels**—Range is 0-1024.<br>• **Concurrent SSL VPN Tunnels**—Range is 0-1024.<br>2. Click **OK**. |
| Step 5 | Save the configuration. | Click **Commit** and **OK**. The virtual system is now an object accessible from the **Objects** tab. |


© Palo Alto Networks, Inc.

| Configure a Virtual System | | |
|---|---|---|
| Step 6 | Create at least one virtual router for the virtual system in order to make the virtual system capable of networking functions, such as static and dynamic routing.<br>Alternatively, your virtual system might use a VLAN or a virtual wire, depending on your deployment. | 1. Select **Network > Virtual Routers** and **Add** a virtual router by **Name**.<br>2. For **Interfaces,** click **Add** and from the drop-down, select the interfaces that belong to the virtual router.<br>3. Click **OK**. |
| Step 7 | Configure a security zone for each interface in the virtual system. | For at least one interface, create a Layer 3 security zone. See Configure Interfaces and Zones. |
| Step 8 | Configure the security policies allowing or denying traffic to and from the zones in the virtual system. | See Set Up Basic Security Policies. |
| Step 9 | Save the configuration. | Click **Commit** and **OK**.<br>After a virtual system is created, a device administrator can use the CLI to commit a configuration for only a specific virtual system:<br>`PA-5060> commit partial vsys vsys<id>` |
| Step 10 | (Optional) View the security policies configured for a virtual system. | Open an SSH session to use the CLI. To view the security policies for a virtual system, in operational mode, use the following commands:<br>`PA-5060> set system setting target-vsys <vsys-id>`<br>`PA-5060> show running security-policy` |

# Configure Inter-Virtual System Communication within the Firewall

Perform this task if you have a use case, perhaps within a single enterprise, where you want the virtual systems to be able to communicate with each other within the firewall. Such a scenario is described in Inter-VSYS Traffic That Remains Within the Firewall. This task presumes:

- ❒ You completed the task, Configure Virtual Systems.
- ❒ When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate with each other to be visible to each other.

| **Configure Inter-Virtual System Communication within the Firewall** | | |
|---|---|---|
| Step 1 | Configure an external zone for each virtual system. | 1. Select **Network > Zones** and **Add** a new zone by **Name**.<br>2. For **Location**, select the virtual system for which you are creating an external zone.<br>3. For **Type**, select **External**.<br>4. For **Virtual Systems**, click **Add** and enter the virtual system that the external zone can reach.<br>5. **Zone Protection Profile**—Optionally select a zone protection profile (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.<br>6. **Log Setting—**Optionally select a log forwarding profile for forwarding zone protection logs to an external system.<br>7. Optionally select the **Enable User Identification** check box to enable User-ID for the external zone.<br>8. Click **OK**. |
| Step 2 | Configure the security policies allowing or denying traffic from the internal zones to the external zone of the virtual system, and vice versa. | • See Set Up Basic Security Policies.<br>• See Inter-VSYS Traffic That Remains Within the Firewall. |
| Step 3 | Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

# Configure a Shared Gateway

Perform this task if you need multiple virtual systems to share an interface (a Shared Gateway) to the Internet. This task presumes:

- You configured an interface with a globally-routable IP address, which will be the shared gateway.
- You completed the prior task, Configure Virtual Systems. For the interface, you chose the external-facing interface with the globally-routable IP address.
- When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate to be visible to each other.

| **Configure a Shared Gateway** | | |
|---|---|---|
| Step 1 | Configure a Shared Gateway. | 1. Select **Device > Shared Gateway**, click **Add** and enter an **ID**.<br>2. Enter a helpful **Name**, preferably including the **ID** of the gateway.<br>3. In the **DNS Proxy** field, select a DNS proxy object if you want to apply DNS proxy rules to the interface.<br>4. **Add** an **Interface** that connects to the outside world.<br>5. Click **OK**. |
| Step 2 | Configure the zone for the shared gateway.<br> When adding objects such as zones or interfaces to a shared gateway, the shared gateway itself will be listed as an available vsys in the **VSYS** drop-down menu. | 1. Select **Network > Zones** and **Add** a new zone by **Name**.<br>2. For **Location**, select the shared gateway for which you are creating a zone.<br>3. For **Type**, select **Layer3**.<br>4. **Zone Protection Profile**—Optionally select a zone protection profile (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.<br>5. **Log Setting—**Optionally select a log forwarding profile for forwarding zone protection logs to an external system.<br>6. Optionally select the **Enable User Identification** check box to enable User-ID for the shared gateway.<br>7. Click **OK**. |
| Step 3 | Save the configuration. | Click **Commit**. |

# Customize Service Routes for a Virtual System


- Customize Service Routes to Services for Virtual Systems
- Configure a PA-7000 Series Firewall for Logging Per Virtual System
- Configure a DNS Proxy Object
- Configure a DNS Server Profile
- Configure Administrative Access Per Virtual System or Device


## Customize Service Routes to Services for Virtual Systems

❒ Prior to performing this task, in order to see the **Global** and **Virtual Systems** tabs, you must enable **Multi Virtual System Capability**.

If **Multi Virtual System Capability** is enabled, any virtual system that does not have specific service routes configured inherits the global service and service route settings for the device.

The firewall supports syslog forwarding on a virtual system basis. When multiple virtual systems on a firewall are connecting to a syslog server using SSL transport, the firewall can generate only one certificate for secure communication. The firewall does not support each virtual system having its own certificate.

In the following use case, you are configuring individual services routes for a firewall with multiple virtual systems.

 © Palo Alto Networks, Inc.

| Customize Service Routes to Services Per Virtual System | |
|---|---|
| Step 1 Customize service routes for a virtual system. | 1. Select **Device > Setup > Services > Virtual Systems**, and select the virtual system you want to configure.<br>2. Click the **Service Route Configuration** link.<br>3. Select one of the radio buttons:<br>• **Inherit Global Service Route Configuration**—Causes the virtual system to inherit the global service route settings relevant to a virtual system. If you choose this option, skip down to step 7.<br>• **Customize**—Allows you to specify a source interface and source address for each service.<br>4. If you chose **Customize**, select the **IPv4** or **IPv6** tab, depending on what type of addressing the server offering the service uses. You can specify both IPv4 and IPv6 addresses for a service. Click the check box(es) for the services for which you want to specify the same source information. (Only services that are relevant to a virtual system are available.) Click **Set Selected Service Routes**.<br>• For **Source Interface**, select **Any**, **Inherit Global Setting**, or an interface from the drop-down to specify the source interface that will be used in packets sent to the external service(s). Hence, the server's response will be sent to that source interface. In our example deployment, you would set the source interface to be the subinterface of the tenant.<br>• **Source Address** will indicate **Inherited** if you selected **Inherit Global Setting** for the **Source Interface** or it will indicate the source address of the **Source Interface** you selected. If you selected **Any** for **Source Interface**, select an IP address from the drop-down, or enter an IP address (using the IPv4 or IPv6 format that matches the tab you chose) to specify the source address that will be used in packets sent to the external service.<br>• If you modify an address object and the IP family type (IPv4/IPv6) changes, a **Commit** is required to update the service route family to use.<br>5. Click **OK**.<br>6. Repeat steps 4 and 5 to configure source addresses for other external services.<br>7. Click **OK**. |
| Step 2 Save the configuration. | Click **Commit** and **OK**.<br>If you are configuring per-virtual system service routes for logging services for a PA-7000 Series firewall, continue to the task Configure a PA-7000 Series Firewall for Logging Per Virtual System. |

## Configure a PA-7000 Series Firewall for Logging Per Virtual System

❒ You must have enabled **Multi Virtual System Capability** (**Device > Setup > Management**) in order to access the LPC subinterface configuration.

Perform this task on your PA-7000 Series firewall to configure logging for different virtual systems. For more information, see PA-7000 Series Firewall LPC Support for Per-Virtual System Paths to Logging Servers.

| **Configure a PA-7000 Series Firewall Subinterface for Service Routes per Virtual System** | | |
|---|---|---|
| Step 1 | Create a Log Card subinterface. | 1. Select **Network > Interfaces > Ethernet** and select the interface that will be the Log Card interface.<br>2. Enter the **Interface Name**.<br>3. For **Interface Type**, select **Log Card** from the drop-down.<br>4. Click **OK**. |
| Step 2 | Add a subinterface for each tenant on the LPCs physical interface. | 1. Highlight the Ethernet interface that is a Log Card interface type and click **Add Subinterface**.<br>2. For **Interface Name**, after the period, enter the subinterface assigned to the tenant's virtual system.<br>3. For **Tag**, enter a VLAN tag value.<br> Make the tag the same as the subinterface number for ease of use, but it could be a different number.<br>4. (Optional) Enter a **Comment**.<br>5. On the **Config** tab, in the **Assign Interface to Virtual System** field, select the virtual system to which the LPC subinterface is assigned (from the drop-down). Alternatively, you can click **Virtual Systems** to add a new virtual system.<br>6. Click **OK**. |
| Step 3 | Enter the addresses assigned to the subinterface, and configure the default gateway. | 1. Select the **Log Card Forwarding** tab, and do one or both of the following:<br>• For the IPv4 section, enter the **IP Address** and **Netmask** assigned to the subinterface. Enter the **Default Gateway** (the next hop where packets will be sent that have no known next hop address in the Routing Information Base [RIB]).<br>• For the IPv6 section, enter the **IPv6 Address** assigned to the subinterface. Enter the **IPv6 Default Gateway**.<br>2. Click **OK**. |
| Step 4 | Save the configuration. | Click **OK** and **Commit**. |
| Step 5 | If you haven't already done so, configure the remaining service routes for the virtual system. | Customize Service Routes for a Virtual System. |

 © Palo Alto Networks, Inc.

## Configure a DNS Proxy Object

If your firewall is to act as a DNS proxy for a virtual system, perform this task to configure a DNS Proxy Object. The proxy object can either be shared among all virtual systems or applied to a specific virtual system.

| Configure a DNS Proxy Object | | |
|---|---|---|
| Step 1 | Configure the basic settings for a DNS Proxy object. | 1. Select **Network > DNS Proxy** and **Add** a new object.<br>2. Verify that **Enable** is selected.<br>3. Enter a **Name** for the object.<br>4. For **Location**, select the virtual system to which the object applies. If you select **Shared**, you must specify at least a **Primary** DNS server address., and optionally a **Secondary** address.<br>5. If you selected a virtual system, for **Server Profile**, select a DNS Server profile or else click **DNS Server Profile** to configure a new profile. see Configure a DNS Server Profile.<br>6. For **Interface**, click **Add** and specify the interfaces to which the DNS Proxy object applies.<br>• If you use the DNS Proxy object for performing DNS lookups, an interface is required. The firewall will listen for DNS requests on this interface, and then proxy them.<br>• If you use the DNS Proxy object for a service route, the interface is optional. |
| Step 2 | (Optional) Specify DNS Proxy rules. | 1. On the **DNS Proxy Rules** tab, click **Add** and enter a **Name** for the rule.<br>2. **Turn on caching of domains resolved by this mapping** if you want the firewall to cache the resolved domains.<br>3. For **Domain Name**, click **Add** and enter one or more domains, one entry per row. Each domain name can contain * as a wildcard. The number of tokens in a wildcard string must match the number of tokens in the requested domain. For example, '*.engineering.local' will not match 'engineering.local'. Both entries must be specified if you want both.<br>4. In Step 4 above, for **Location**:<br>• If you chose a virtual system, select a **DNS Server profile** here.<br>• If you chose **Shared**, enter a **Primary** address here.<br>5. Click **OK**. |
| Step 3 | (Optional) Supply the DNS Proxy with static FQDN-to-address entries. Static DNS entries allow the firewall to resolve the FQDN to an IP address without going out to the DNS server. | 1. On the **Static Entries** tab, click **Add** and enter a **Name**.<br>2. Enter the Fully Qualified Domain Name (**FQDN**).<br>3. For **Address**, click **Add** and enter the IP address to which the FQDN should be mapped.<br>4. Repeat steps 1-3 to provide additional static entries.<br>5. Click **OK**. |

© Palo Alto Networks, Inc.

| Configure a DNS Proxy Object (Continued) | | |
|---|---|---|
| Step 4 | (Optional) Enable caching and configure other advanced settings for the DNS Proxy. | 1. On the **Advanced** tab, click **Cache** to enable the firewall to cache FQDN-to-address mappings that the firewall learns.<br>• **Size**—Enter the maximum number of entries the firewall can cache (range is 1024-10240; default is 1024).<br>• **Timeout**—Enter the number of hours after which all cached entries are removed (range is 4-24; default is 4). DNS time-to-live values are used to remove cache entries when they have been stored for less than the configured timeout period. After a timeout, new DNS requests must be resolved and cached again.<br>2. Select **TCP Queries** to enable DNS queries using TCP.<br>• **Max Pending Requests**—Enter the maximum number of concurrent, pending TCP DNS requests that the firewall will support (range is 64-256; default is 64).<br>3. For **UDP Queries Retries**, enter the following:<br>• **Interval**—Enter the length of time (in seconds) after which another request is sent if no response has been received. (range is 1-30; default is 2).<br>• **Attempts**—Enter the maximum number of UDP query attempts (excluding the first attempt) after which the next DNS server is queried (range is 1-30; default is 5.) |
| Step 5 | Save the configuration. | Click **OK** and **Commit**. |

## Configure a DNS Server Profile

Perform this task to configure a DNS Server Profile, which simplifies configuration of a virtual system. The **Primary DNS** or **Secondary DNS** address is used to create the DNS request that the virtual system sends to the DNS server.

 © Palo Alto Networks, Inc.

| Configure a DNS Server Profile | | |
|---|---|---|
| Step 1 | Name the DNS server profile, select the virtual system to which it applies, and specify the primary and secondary DNS server addresses. | 1. Select **Device > Server Profiles > DNS** and click **Add**.<br>2. Enter a **Name** for the DNS server profile.<br>3. For **Location**, select the virtual system to which the profile applies.<br>4. For **Inheritance Source**, from the drop-down, select **None** if the DNS server addresses are not inherited. Otherwise, specify the DNS server from which the profile should inherit settings. If you choose a DNS server, click **Check inheritance source status** to see that information.<br>5. Specify the IP address of the **Primary DNS** server, or leave as **inherited** if you chose an **Inheritance Source**.<br> Keep in mind that if you specify an FQDN instead of an IP address, the DNS for that FQDN is resolved in **Device > Virtual Systems > DNS Proxy**.<br>6. Specify the IP address of the **Secondary DNS** server, or leave as **inherited** if you chose an **Inheritance Source**. |
| Step 2 | Configure the service route that the firewall automatically uses, based on whether the target DNS Server has an IP address family type of IPv4 or IPv6. | 1. Click **Service Route IPv4** to enable the subsequent interface and IPv4 address to be used as the service route, if the target DNS address is an IPv4 address.<br>2. Specify the **Source Interface** to select the DNS server's source IP address that the service route will use. The firewall determines which virtual router is assigned that interface, and then does a route lookup in the virtual router routing table to reach the destination network (based on the **Primary DNS** address).<br>3. Specify the IPv4 **Source Address** from which packets going to the DNS server are sourced.<br>4. Click **Service Route IPv6** to enable the subsequent interface and IPv6 address to be used as the service route, if the target DNS address is an IPv6 address.<br>5. Specify the **Source Interface** to select the DNS server's source IP address that the service route will use. The firewall determines which virtual router is assigned that interface, and then does a route lookup in the virtual router routing table to reach the destination network (based on the **Primary DNS** address).<br>6. Specify the IPv6 **Source Address** from which packets going to the DNS server are sourced.<br>7. Click **OK**. |
| Step 3 | Save the configuration. | Click **OK** and **Commit**. |

## Configure Administrative Access Per Virtual System or Device

If you have a superuser administrative account, you now have the ability to create and configure more granular permissions for a vsysadmin or device admin role.

| Create an Admin Role Profile Per Virtual System or Device | | |
|---|---|---|
| Step 1 | Create an Admin Role Profile that grants or disables permission to an Administrator to configure or read-only various areas of the web interface. | 1. Select **Device > Admin Roles** and **Add** an **Admin Role Profile**.<br>2. Enter a **Name** and optional **Description** of the profile.<br>3. For **Role**, specify which level of control the profile affects:<br>• **Device**—The profile allows the management of the global settings and any virtual systems.<br>• **Virtual System**—The profile allows the management of only the virtual system(s) assigned to the administrator(s) who have this profile. (The administrator will be able to access **Device > Setup > Services > Virtual Systems**, but not the **Global** tab.)<br>4. On the **Web UI** tab for the Admin Role Profile, scroll down to **Device**, and leave the green check mark (Enable).<br>• Under **Device**, enable **Setup**. Under **Setup**, enable the areas to which this profile will grant configuration permission to the administrator, as shown below. (The Read Only lock icon appears in the Enable/Disable rotation if Read Only is allowed for that setting.)<br>– **Management**—Allows an admin with this profile to configure settings on the **Management** tab.<br>– **Operations**—Allows an admin with this profile to configure settings on the **Operations** tab.<br>– **Services**—Allows an admin with this profile to configure settings on the **Services** tab. An admin must have **Services** enabled in order to access the **Device > Setup Services > Virtual Systems** tab. If the **Role** was specified as **Virtual System** in the prior step, **Services** is the only setting that can be enabled under **Device > Setup**.<br>– **Content-ID**—Allows an admin with this profile to configure settings on the **Content-ID** tab.<br>– **WildFire**—Allows an admin with this profile to configure settings on the **WildFire** tab.<br>– **Session**—Allows an admin with this profile to configure settings on the **Session** tab.<br>– **HSM**—Allows an admin with this profile to configure settings on the **HSM** tab.<br>5. Click **OK**.<br>6. (Optional) Repeat the entire step to create another Admin Role profile with different permissions, as necessary. |

 © Palo Alto Networks, Inc.

| Create an Admin Role Profile Per Virtual System or Device (Continued) | | |
|---|---|---|
| Step 2 | Apply the Admin role profile to an administrator. | 1. Select **Device > Administrators**, click **Add** and enter the **Name** to add an Administrator.<br>2. (Optional) Select an **Authentication Profile**.<br>3. (Optional) Select **Use only client certificate authentication (Web)** to have bi-directional authentication; to get the server to authenticate the client.<br>4. Enter a **Password** and **Confirm Password**.<br>5. (Optional) Select **Use Public Key Authentication (SSH)** if you want to use a much stronger, key-based authentication method using an SSH public key rather than just a password.<br>6. For **Administrator Type**, select **Role Based**.<br>7. For **Profile**, select the profile that you just created.<br>8. (Optional) Select a **Password Profile**.<br>9. Click **OK**. |
| Step 3 | Save the configuration. | Click **Commit** and **OK**. |

# DNS Resolution—Three Use Cases

The firewall determines how to handle DNS requests based on where the request originated. This section illustrates three types of DNS resolution, which are listed in the following table. The binding location determines which DNS proxy object is used for the resolution. For illustration purposes, the use cases show how a service provider might configure DNS settings to provide DNS services for resolving DNS queries required on the firewall and for tenant (subscriber) virtual systems.

| Resolution Type | Location: Shared | Location: Specific Vsys |
|---|---|---|
| Firewall DNS resolution—performed by management plane | Binding: Global<br>Illustrated in Use Case 1 | N/A |
| Security profile, reporting, and server profile resolution—performed by management plane | Binding: Global<br>Same behavior as Use Case 1 | Binding: Specific vsys<br>Illustrated in Use Case 2 |
| DNS proxy resolution for DNS client hosts connected to interface on firewall, going through the firewall to a DNS Server—performed by dataplane | Binding: Interface<br>Service Route: Interface and IP address on which the DNS Request was received.<br>Illustrated in Use Case 3 | |

- Use Case 1: Firewall Requires DNS Resolution for Management Purposes
- Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System
- Use Case 3: Firewall Acts as DNS Proxy Between Client and Server

## Use Case 1: Firewall Requires DNS Resolution for Management Purposes

In this use case, the firewall is the client requesting DNS resolutions of FQDNs for management events such as software update services, dynamic software updates, or WildFire. The shared, global DNS services perform the DNS resolution for the management plane functions.

 © Palo Alto Networks, Inc.

| Configure DNS Services for the Firewall | |
|---|---|
| Step 1 Configure the primary and secondary DNS servers you want the firewall to use for its management DNS resolutions.<br><br>You must manually configure at least one DNS server on the firewall or it will not be able to resolve hostnames; it will not use DNS server settings from another source, such as an ISP. | 1. Select **Device > Setup > Services > Global** and Edit. (For devices that do not support multiple virtual systems, there is no **Global** tab; simply edit the Services.)<br>2. On the **Services** tab, for **DNS**, click **Servers** and enter the **Primary DNS Server** address and **Secondary DNS Server** address.<br> <br>3. Click **OK** and **Commit**. |

**Configure DNS Services for the Firewall (Continued)**

Step 2 Alternatively, you can configure a DNS Proxy Object if you want to configure advanced DNS functions such as split DNS, DNS proxy overrides, DNS proxy rules, static entries, or DNS inheritance.

1. On the **Services** tab, for **DNS**, click **DNS Proxy Object**.

2. From the **DNS Proxy** drop-down, select the DNS proxy that you want to use to configure global DNS services, or click **DNS Proxy** to configure a new DNS proxy object, as shown in the following screenshot and subsequent steps.

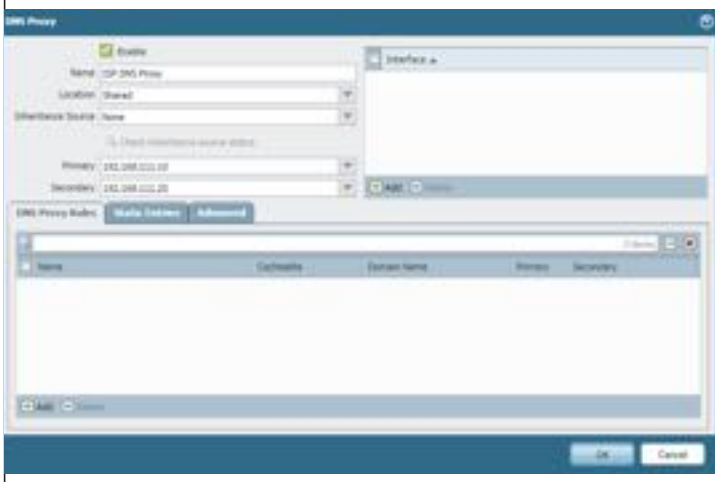

3. To create a new proxy object, click **Enable** and enter a **Name** for the DNS proxy object.
4. For **Location**, select **Shared** for global, device-wide DNS proxy services.

   Shared DNS proxy objects do not use DNS server profiles because they do not require a specific service route belonging to a tenant virtual system.
5. For **Primary**, enter the primary DNS server IP address. Optionally enter a **Secondary** DNS server IP address. In the ISP example in the screenshot above, the DNS proxy defines the primary and secondary DNS servers that are used to resolve the firewall management services.
6. Click **OK** and **Commit**.


© Palo Alto Networks, Inc.

## Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System

In this use case, multiple tenants (ISP subscribers) are defined on the firewall and each tenant is allocated a separate virtual system (vsys) and virtual router in order to segment its services and administrative domains. The following figure illustrates several virtual systems within a firewall.

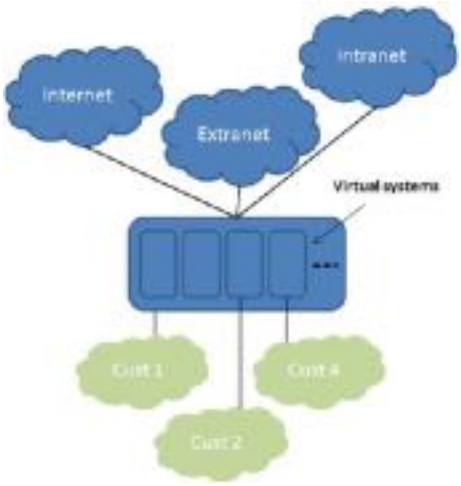


Each tenant has its own server profiles for its security policies, reporting, and management services (such as email, Kerberos, SNMP, syslog, and more) defined in its own networks.

For the DNS resolutions initiated by these services, each virtual system is configured with its own DNS Proxy object to allow each tenant to customize how DNS resolution is handled within its virtual system. Any service with a **Location** will use the DNS Proxy object configured for the virtual system to determine the primary (or secondary) DNS server to resolve FQDNs, as illustrated in the following figure.

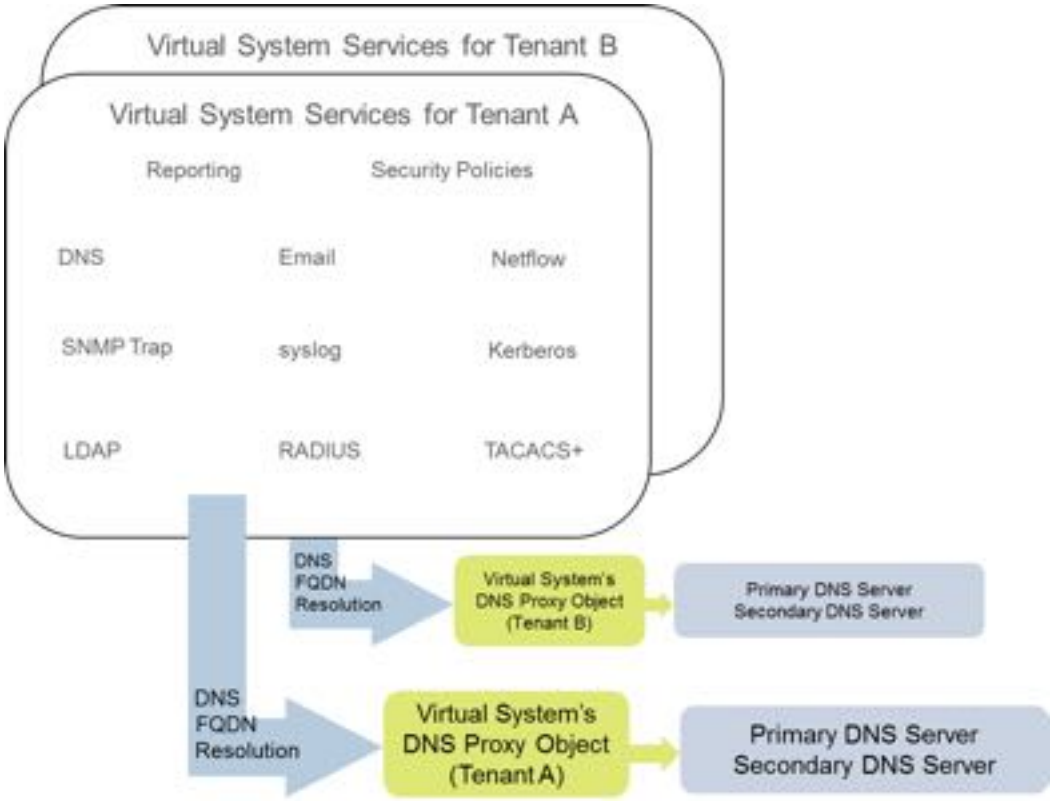


 © Palo Alto Networks, Inc.

| Configure a DNS Proxy for a Virtual System | |
|---|---|
| Step 1 For each virtual system, specify the DNS Proxy to use. | 1. Select **Device > Virtual Systems** and click **Add**.<br>2. Enter the **ID** of the virtual system (range is 1-255), and an optional **Name**, in this example, Corp1 Corporation.<br>3. On the **General** tab, choose a **DNS Proxy** or create a new one. In this example, Corp1 DNS Proxy is selected as the proxy for Corp1 Corporation's virtual system.<br>(If you need to create a new DNS Proxy, Step 2 below shows how to create a DNS Proxy and a Server Profile.)<br>4. For **Interfaces**, click **Add**. In this example, Ethernet1/20 is dedicated to this tenant.<br>5. For **Virtual Routers**, click **Add**. A virtual router named Corp1 VR is assigned to the virtual system in order to separate routing functions.<br>6. Click **OK** to save the configuration.<br> 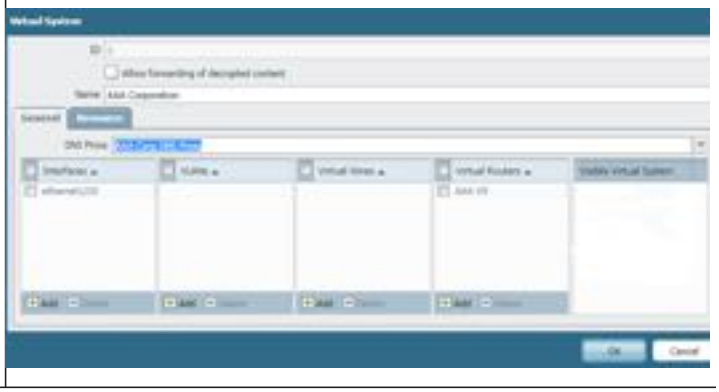 |

**Configure a DNS Proxy for a Virtual System**

Step 2 Configure a DNS Proxy and a server profile to support DNS resolution for a virtual system.

1. Select **Network > DNS Proxy** and click **Add**.
2. Click **Enable** and enter a **Name** for the DNS Proxy.
3. For **Location**, select the virtual system of the tenant, in this example, Corp1 Corporation (vsys6). (You could choose the **Shared** DNS Proxy resource instead.)

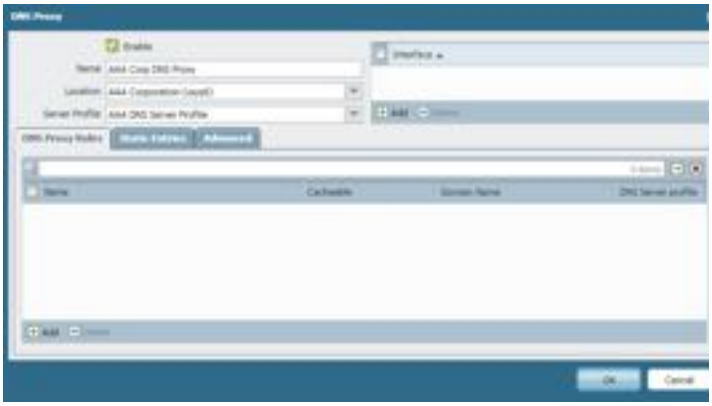

4. For **Server Profile**, choose or create a profile to customize DNS servers to use for DNS resolutions for this tenant's security policy, reporting, and server profile services.

   If the profile is not already configured, in the **Server Profile** field, click **DNS Server Profile** to Configure a DNS Server Profile. In this example, the Corp1 DNS Server Profile was created.

   The DNS server profile identifies the IP addresses of the primary and secondary DNS server to use for management DNS resolutions for this virtual system.
5. Also for this server profile, optionally configure a **Service Route IPv4** and/or a **Service Route IPv6** to instruct the firewall which **Source Interface** to use in its DNS requests. If that interface has more than one IP address, configure the **Source Address** also.
6. Click **OK** to save the DNS Server Profile.
7. Click **OK** and **Commit** to save the DNS Proxy.

 © Palo Alto Networks, Inc.

Optional advanced features such as split DNS can be configured using **DNS Proxy Rules**. A separate DNS server profile can be used to redirect DNS resolutions matching the **Domain Name** in a **DNS Proxy Rule** to another set of DNS servers, if required. Use Case 3 illustrates split DNS.

If you use two separate DNS server profiles in the same DNS Proxy object, one for the DNS Proxy and one for the DNS proxy rule, the following behaviors occur:

- If a service route is defined in the DNS server profile used by the DNS Proxy, it takes precedence and is used.
- If a service route is defined in the DNS server profile used in the DNS proxy rules, it is not used. If the service route differs from the one defined in the DNS server profile used by the DNS Proxy, the following warning message is displayed during the **Commit** process:

  ```
  Warning: The DNS service route defined in the DNS proxy object is different from the DNS proxy rule's service route. Using the DNS proxy object's service route.
  ```

- If no service route is defined in any DNS server profile, the global service route is used if needed.

## Use Case 3: Firewall Acts as DNS Proxy Between Client and Server

In this use case, the firewall is located between a DNS client and a DNS server. A DNS Proxy on the firewall is configured to act as the DNS server for the hosts that reside on the tenant's network connected to the firewall interface. In such a scenario, the firewall performs DNS resolution on its dataplane.

This scenario happens to use *split DNS*, a configuration where **DNS Proxy Rules** are configured to redirect DNS requests to a set of DNS servers based on a domain name match. If there is no match, the **Server Profile** determines the DNS servers to which the request is sent, hence the two, split DNS resolution methods.

For dataplane DNS resolutions, the source IP address from the DNS proxy in PAN-OS to the outside DNS server would be the address of the proxy (the destination IP of the original request). Any service routes defined in the DNS Server Profile are not used. For example, if the request is from host 1.1.1.1 to the DNS proxy at 2.2.2.2, then the request to the DNS server (at 3.3.3.3) would use a source of 2.2.2.2 and a destination of 3.3.3.3.

| Configure a DNS Proxy and DNS Proxy Rules | |
|---|---|
| Step 1 Configure a DNS Proxy and DNS proxy rules. | 1. Select **Network > DNS Proxy** and click **Add**.<br>2. Click **Enable** and enter a **Name** for the DNS Proxy.<br>3. For **Location**, select the virtual system of the tenant, in this example, Corp1 Corporation (vsys6).<br>4. For **Interface**, select the interface that will receive the DNS requests from the tenant's hosts, in this example, Ethernet1/20.<br>5. Choose or create a **Server Profile** to customize DNS servers to resolve DNS requests for this tenant. In this example, the Corp1 DNS Server Profile is selected.<br> 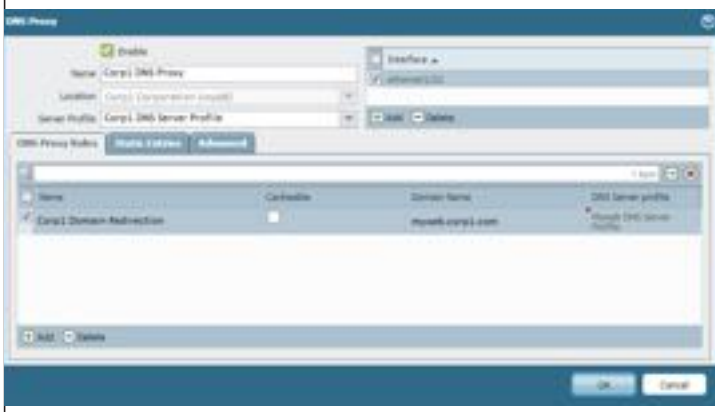 <br>6. On the **DNS Proxy Rules** tab, click **Add** and enter a **Name** for the rule.<br>7. Optionally select **Turn on caching of domains resolved by this mapping**.<br>8. Click **Add** and enter one or more **Domain Name**(s), one entry per row.<br>Each domain name can contain * as a wild card. The number of characters in a wildcard string must equal the number of characters in the requested domain to match. For example, '*.engineering.local' does not match 'engineering.local'. Both domain names must be specified in order for both to be matched.<br>9. For **DNS Server profile**, select a profile from the drop-down. The firewall compares the domain name in the DNS request to the domain name(s) defined in the **DNS Proxy Rules**. If there is a match, the **DNS Server profile** defined in the rule is used to determine the DNS server.<br>In this example, if the domain in the request matches myweb.corp1.com, the DNS server defined in the myweb DNS Server Profile is used. If there is no match, the DNS server defined in the **Server Profile** (Corp1 DNS Server Profile) is used.<br>10. Click **OK** to save the rule.<br>11. Click **OK** to save the DNS Proxy. |

 © Palo Alto Networks, Inc.

# Virtual System Functionality with Other Features

Many of the firewall's features and functionality are capable of being configured, viewed, logged, or reported per virtual system. Therefore, virtual systems are mentioned in other relevant locations in the documentation and that information is not repeated here. Some of the specific chapters are the following:

- If you are configuring Active/Passive HA, the two firewalls must have the same virtual system capability (single or multiple virtual system capability). See High Availability.
- To configure QoS for virtual systems, see Configure QoS for a Virtual System.
- For information about configuring a firewall with virtual systems in a virtual wire deployment that uses subinterfaces (and VLAN tags), see the Virtual Wire Subinterfaces in Interface Deployments.

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

# Certifications

The following topics describe how to configure the firewall to support the Common Criteria and the Federal Information Processing Standards 140-2 (FIPS 140-2), which are security certifications that ensure a standard set of security assurances and functionalities. These certifications are often required by civilian U.S. government agencies and government contractors.

- Enable FIPS and Common Criteria Support
- CCEAL4 Security Functions

# Enable FIPS and Common Criteria Support

Use the following procedure to enable CCEAL4 mode on a software version that supports Common Criteria and the Federal Information Processing Standards 140-2 (FIPS 140-2). When you enable CCEAL4 mode, all FIPS functionality is also included in this mode.

When you enable CCEAL4 mode, the device will reset to the factory default settings; all configuration will be removed.

| Enable CCEAL4 Mode | |
|---|---|
| Step 1 | Boot the device into maintenance mode as follows:<br>1. Establish a serial connection to the console port on the firewall.<br>2. Enter the following CLI command:<br>`debug system maintenance-mode`<br>3. Press Enter to continue.<br>You can also reboot the firewall and enter `maint` at the maintenance mode prompt. |
| Step 2 | Select **Set CCEAL4 Mode** from the menu.<br>Official certification support is only provided with **CCEAL4** mode; not **FIPS** mode. |
| Step 3 | Select **Enable CCEAL4 Mode** from the menu. |
| Step 4 | When prompted, select **Reboot**.<br>After successfully switching to CCEAL4 mode, the following status displays: `CCEAL4 mode enabled successfully`. In addition, the following changes will take place:<br>• **CC** will display at all times in the status bar at the bottom of the web interface.<br>• The console port functions as a status output port only<br>• The default admin login credentials change to admin/paloalto. |

 © Palo Alto Networks, Inc.

# CCEAL4 Security Functions

When CCEAL4 mode is enabled, the following security functions are enforced:

- To log into the firewall, the browser must be TLS 1.0 (or later) compatible. On a WF-500 appliance, you manage the appliance using the CLI only and you must connect using an SSHv2 compatible client application.
- All passwords on the firewall must be at least six characters.
- You must enforce a **Failed Attempts** and **Lockout Time (min)** value that is greater than 0 in authentication settings. If an administrator reaches the **Failed Attempts** threshold, the administrator is locked out for the duration defined in the **Lockout Time (min)** field.
- You must enforce an **Idle Timeout** value greater than 0 in authentication settings. If a login session is idle for more than the specified value, the account is automatically logged out.
- The firewall automatically determines the appropriate level of self-testing and enforces the appropriate level of strength in encryption algorithms and cipher suites.
- Non-CC/FIPS approved algorithms are not decrypted and are thus ignored during decryption.
- When configuring an IPSec VPN, the administrator must select a cipher suite option presented to them during the IPSec setup.
- Self-generated and imported certificates must contain public keys that are either RSA 2048 bits (or more) or ECDSA 256 bits (or more) and you must use a digest of SHA256 or greater.
- The serial console port is only available as a status output port when CCEAL4 is enabled.
- Telnet, TFTP, and HTTP management connections are unavailable.
- High availability (HA) port encryption is required.

© Palo Alto Networks, Inc.

 © Palo Alto Networks, Inc.

paloalto NETWORKS®

# PAN-OS® New Features Guide

Version 7.0

## Contact Information

Corporate Headquarters:

Palo Alto Networks

4401 Great America Parkway

Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-us

## About this Guide

This guide describes how to use the new features introduced in PAN-OS 7.0. For additional information, refer to the following resources:

- For information on the additional capabilities and for instructions on configuring the features on the firewall, refer to https://www.paloaltonetworks.com/documentation.
- For access to the knowledge base and community forums, refer to https://live.paloaltonetworks.com.
- For contacting support, for information on support programs, to manage your account or devices, or to open a support case, refer to https://www.paloaltonetworks.com/support/tabs/overview.html.
- For the most current PAN-OS and Panorama 7.0 release notes, go to https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os-release-notes.html.

To provide feedback on the documentation, please write to us at: documentation@paloaltonetworks.com.

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2016 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at http://www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

**Revision Date: June 8, 2016**

# Authentication Features


- Authentication and Authorization Enhancements
- SSL/TLS Service Profiles
- TACACS+ Authentication
- Kerberos V5 Single Sign-On
- Suite B Cryptography Support
- Authentication Server Connectivity Testing




Copyright © 2007-2015 Palo Alto Networks

# Authentication and Authorization Enhancements

This release provides the following authentication and authorization enhancements:

| Enhancement | Description |
|---|---|
| Easier configuration workflow for authentication servers, profiles, and sequences | The workflow to configure server profiles, authentication profiles, and authentication sequences is simpler and more intuitive. Authentication server profiles now have only the parameters for accessing authentication services (for example, IP addresses), while authentication profiles and sequences have only the parameters for the authentication process (for example, user domain). You no longer need to set the allowed login attempts and account lockout duration for authentication sequences, only authentication profiles.<br>To enable authentication, configure a server profile (LDAP, RADIUS, Kerberos, or TACACS+) if you will use an external authentication service, configure an authentication profile (to which you assign the server profile), optionally assign multiple authentication profiles to an authentication sequence, and assign the authentication profile or sequence to administrator accounts, GlobalProtect portals or gateways, or the Captive Portal configuration. For the relevant procedures, refer to:<br>• Manage firewall administrators<br>• Configure the GlobalProtect Portal<br>• Configure GlobalProtect Gateways<br>• Map IP Addresses to User Names Using Captive Portal<br> You can perform Authentication Server Connectivity Testing to check whether the Palo Alto Networks device can successfully communicate with the authentication server specified in an authentication profile. |
| Kerberos V5 single sign-on (SSO) support | Palo Alto Networks devices now support Kerberos V5 single sign-on, and automatic modifications to user-entered domains and usernames. This reduces the number of logins and the amount of typing required to authenticate administrators and end users. You configure these features in authentication profiles. |
| Certificate verification for LDAP connections | To improve security, you can now enable Palo Alto Network devices to verify the certificate that an LDAP server presents for SSL/TLS connections. You enable the verification in an LDAP server profile. |
| Intuitive configuration workflow for group mapping | The workflow to configure username to group mapping for User-ID is more intuitive: you set the user domain override in the group mapping configuration instead of in LDAP server profiles. |
| Delivery of VSAs from GlobalProtect clients to RADIUS servers | You can now configure the firewall send Vendor-Specific Attributes (VSAs) from GlobalProtect clients to RADIUS servers so that RADIUS administrators can perform administrative tasks based on those VSAs. For example, RADIUS administrators might use the client operating system attribute to define a policy that mandates regular password authentication for Microsoft Windows users and one-time password (OTP) authentication for Google Android users. RADIUS administrators can also log in to the firewall CLI as SSH users without first logging in to the web interface. |
| CLI commands for authentication troubleshooting | The device CLI has new commands for displaying authentication statistics, debugging authentication events, and unlocking users. |



Copyright © 2007-2015 Palo Alto Networks

# SSL/TLS Service Profiles

You can now assign SSL/TLS service profiles to device services that use SSL/TLS, including Captive Portal, GlobalProtect portals and gateways, management traffic access using the web interface or XML API, URL Admin Override, and the User-ID Syslog listening service. SSL/TLS service profiles specify a certificate and the allowed protocol version or range of versions (now including TLSv1.2). By defining the protocol versions, the profiles enable you to restrict the cipher suites that are available to secure communication with the clients requesting the services. This improves network security by enabling devices to avoid SSL/TLS versions that have known weaknesses. If a service request involves a protocol version that is outside the specified range, the device downgrades or upgrades the connection to a supported version.

| Configure and Assign an SSL/TLS Service Profile | | |
|---|---|---|
| Step 1 | Configure the SSL/TLS service profile. | 1. For each desired service, Generate a Certificate.<br>Use only signed certificates for SSL/TLS services, not certificate authority (CA) certificates.<br>2. Select **Device > Certificate Management > SSL/TLS Service Profile**.<br>3. For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available.<br>4. Click **Add** and enter a **Name** to identify the profile.<br>5. Select the **Certificate** you just generated.<br>6. Define the range of protocols that the service can use:<br>• For the **Min Version**, select the earliest TLS version to support: **TLSv1.0**, **TLSv1.1**, or **TLSv1.2**.<br>• For the **Max Version**, select the latest TLS version to support: **TLSv1.0**, **TLSv1.1**, **TLSv1.2**, or **Max** (the latest available version.<br>7. Click **OK** and **Commit**. |
| Step 2 | Assign the profile to the desired service.<br>SSL/TLS service profiles replace certificates for all these services except the User-ID syslog listener, for which the profile is a new setting. | • Content-ID: URL Admin Override—Select **Device > Setup > Content-ID** and **Add** or edit the URL Admin Override entries.<br>• Captive Portal—Select **Device > User Identification > Captive Portal Settings** and edit the settings.<br>• GlobalProtect portals—Select **Network > GlobalProtect > Portals**, **Add** or edit a portal, select the **Portal Configuration** tab, and edit the Network Settings.<br>• GlobalProtect gateways—Select **Network > GlobalProtect > Gateways**, **Add** or edit a gateway, select the **General** tab, and edit the Network Settings.<br>• Management interface: inbound traffic—Select **Device > Setup > Management** and edit the General settings.<br>• User-ID user mapping, syslog listening service—Select **Device > User Identification > User Mapping**, edit the Palo Alto Networks User-ID Agent Setup settings, select the **Server Monitor** tab, and select a **Syslog Service Profile**. The SSL/TLS service profile secures syslog messages that the firewall parses when it performs user mapping. |



Copyright © 2007-2015 Palo Alto Networks

# TACACS+ Authentication

Palo Alto Networks devices now support Terminal Access Controller Access-Control System Plus (TACACS+) protocol for authenticating administrator and end users. TACACS+ provides greater security than RADIUS because it encrypts usernames and passwords (instead of just passwords), and is also more reliable (it uses TCP instead of UDP).

| Configure TACACS+ Authentication | | |
|---|---|---|
| Step 1 | Configure a TACACS+ server profile. | 1. Select **Device > Server Profiles > TACACS+** and click **Add**.<br>2. Enter a **Profile Name** to identify the server profile.<br>3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the **Location** where the profile is available.<br>4. Enter a **Timeout** interval after which an authentication request times out (1-20 seconds, default 3 seconds).<br>5. Select the **Use single connection for all authentication** check box to use the same TCP session for all authentications that use this profile. This option improves performance by avoiding the need to start and end a separate TCP session for each authentication. The check box is cleared by default.<br>6. For each TACACS+ server, click **Add** and enter a **Name** (to identify the server), server IP address or FQDN (**TACACS+ Server** field), **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords), and server **Port** for authentication requests (default 49).<br>7. Click **OK** and **Commit**. |
| Step 2 | Use the profile for authentication. | Assign the TACACS+ server profile to an authentication profile and assign the authentication profile to administrator accounts, GlobalProtect portals or gateways, or the Captive Portal configuration. For the specific steps, refer to:<br>• Create an administrative account<br>• Configure the GlobalProtect Portal<br>• Configure GlobalProtect Gateways<br>• Map IP Addresses to User Names Using Captive Portal |



Copyright © 2007-2015 Palo Alto Networks

# Kerberos V5 Single Sign-On

Palo Alto Networks devices now support Kerberos V5 single sign-on (SSO) for authenticating administrative users to the web interface and end users to Captive Portal. A network that supports Kerberos SSO prompts a user to log in only for initial access to the network (for example, logging in to Microsoft Windows). After this initial login, the user can access any browser-based service in the network (for example, the firewall web interface) without having to log in again, until the SSO session expires. (Your Kerberos administrator sets the duration of SSO sessions.)

You enable SSO for a Palo Alto Networks device by importing a Kerberos *keytab* into an authentication profile and assigning the profile to administrator accounts or the Captive Portal configuration. A keytab is a file that contains Kerberos account information (principal name and hashed password) for the device, which is required for SSO authentication. Each authentication profile can have one keytab. If SSO authentication fails, the device prompts the user to log in manually and performs authentication of the type specified in the profile (for example, RADIUS).

- Configure Kerberos SSO for Administrator Authentication
- Configure Kerberos SSO for Captive Portal Authentication

## Configure Kerberos SSO for Administrator Authentication

Configuring Kerberos SSO for administrator authentication has the following prerequisites:

- Your network has a Kerberos infrastructure set up, including a Key Distribution Center (KDC) with an Authentication Server (AS) and Ticket Granting Service (TGS).
- The Palo Alto Networks device that will authenticate users has a Kerberos account, including a principal name and password. You will use these to create the keytab.



Copyright © 2007-2015 Palo Alto Networks

| Configure Kerberos SSO for Administrator Authentication | | |
|---|---|---|
| Step 1 | Create a Kerberos keytab for the Palo Alto Networks device.<br> <br>If the device is in FIPS or CC mode, the algorithm must be `aes128-cts-hmac-sha1-96` or `aes256-cts-hmac-sha1-96`. Otherwise, you can also use `des3-cbc-sha1` or `arcfour-hmac`. To use an AES algorithm, the functional level of the KDC must be Windows Server 2008 or later and you must enable AES encryption for the device account.<br>The algorithm in the keytab must match the algorithm in the service ticket that the TGS issues to clients. Otherwise, the SSO process will fail. Your Kerberos administrator determines which algorithms the service tickets use. | 1. Log in to the KDC and open a command prompt.<br>2. Enter the following command, where <principal_name>, <password>, and <algorithm> are variables. The Kerberos principal name and password are of the device, not the user.<br>**ktpass /princ** `<principal_name>` **/pass** `<password>` **/crypto** `<algorithm>` **/ptype KRB5_NT_PRINCIPAL /out** `<file_name>`**.keytab**<br>For example:<br>**/princ host/computer.abrealm.com@ABREALM.COM /pass securepwd /crypto des3-cbc-sha1 /ptype KRB5_NT_PRINCIPAL /out device.keytab** |
| Step 2 | (External authentication only) Configure an authentication server profile to authenticate administrators if the SSO process fails. | • Configure a RADIUS Server Profile.<br>• Configure an LDAP Server Profile.<br>• Configure a Kerberos Server Profile.<br>• Configure a TACACS+ server profile. |
| Step 3 | Configure an authentication profile. | 1. Select **Device > Authentication Profile**, click **Add**, and enter a **Name** to identify the authentication profile.<br>2. (External authentication only) If the authentication **Type** is an external server, assign the **Server Profile** you just created.<br>3. (New) Enter the **Kerberos Realm** and import the **Kerberos Keytab** you just created. Usually the realm is the DNS domain name of the administrator, except that the realm name is uppercase.<br>4. In the **Advanced** tab, **Add** the administrator user, user group, or **all** to the **Allow List**, and then click **OK**. |
| Step 4 | Assign the authentication profile to the administrator account. | 1. Configure the administrator account.<br>2. Click **Commit**. |



Copyright © 2007-2015 Palo Alto Networks

# Configure Kerberos SSO for Captive Portal Authentication

Configuring Kerberos SSO for Captive Portal authentication has the following prerequisites:

- Your network has a Kerberos infrastructure set up, including a Key Distribution Center (KDC) with an Authentication Server (AS) and Ticket Granting Service (TGS).
- The Palo Alto Networks device that will authenticate users has a Kerberos account, including a principal name and password. You will use these to create the keytab.
- You have set up the initial Captive Portal and policy configurations. In the following procedure, you modify the configurations by configuring redirect mode.

If you configure both Kerberos SSO and NT LAN Manager (NTLM) authentication, the firewall tries Kerberos SSO authentication first, and if that fails, falls back to NTLM authentication for Microsoft Windows clients. If NTLM authentication also fails, the firewall falls back to web form or client certificate authentication, depending on your Captive Portal configuration.

| Configure Kerberos SSO for Captive Portal Authentication | |
|---|---|
| Step 1 Create a Kerberos keytab for the redirect host you will assign to the Captive Portal configuration. The redirect host is the intranet hostname that resolves to the IP address of the firewall Layer 3 interface to which you are redirecting requests that match a Captive Portal policy.<br> <br>If the device is in FIPS or CC mode, the algorithm must be `aes128-cts-hmac-sha1-96` or `aes256-cts-hmac-sha1-96`. Otherwise, you can also use `des3-cbc-sha1` or `arcfour-hmac`. To use an AES algorithm, the functional level of the KDC must be Windows Server 2008 or later and you must enable AES encryption for the device account.<br>The algorithm in the keytab must match the algorithm in the service ticket that the TGS issues to clients to enable single sign-on (SSO). Otherwise, the SSO process will fail. Your Kerberos administrator determines which algorithms the service tickets use. | 1. Log in to the KDC and open a command prompt.<br>2. Enter the following command, where <principal_name>, <password>, and <algorithm> are variables. The Kerberos principal name and password are of the device, not the user.<br>**ktpass /princ** `<principal_name>` **/pass** `<password>` **/crypto** `<algorithm>` **/ptype KRB5_NT_PRINCIPAL /out** `<file_name>`**.keytab**<br>For example:<br>**/princ host/computer.abrealm.com@ABREALM.COM /pass securepwd /crypto des3-cbc-sha1 /ptype KRB5_NT_PRINCIPAL /out device.keytab** |



Copyright © 2007-2015 Palo Alto Networks

| Configure Kerberos SSO for Captive Portal Authentication (Continued) | | |
|---|---|---|
| Step 2 | Enable Captive Portal redirect requests on a Layer 3 interface. | 1. Create an Interface Management profile so the interface can display Captive Portal response pages:<br>a. Select **Network > Network Profiles > Interface Mgmt** and click **Add**.<br>b. Enter a **Name** for the profile, select **Response Pages**, and click **OK**.<br>2. Assign the Interface Management profile to a Layer 3 interface:<br>a. Select **Network > Interfaces > Ethernet** and click the Name of a Layer 3 interface.<br>b. Select **Advanced > Other**, select the **Management Profile** you just created, and click **OK**.<br>3. Create DNS address (A) and pointer (PTR) records that map the IP address on the Layer 3 interface to the redirect host. |
| Step 3 | To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface in Step 2. You can generate and import a self-signed certificate (as described in this example) or import a certificate that an external certificate authority (CA) signed. | To use a self-signed certificate, create a root CA certificate and use it to sign the certificate you will use for Captive Portal:<br>1. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Generate a root CA certificate. Be sure to select the **Certificate Authority** check box.<br>3. Generate a certificate to use for Captive Portal. Be sure to configure the following fields:<br>• **Common Name**—Enter the DNS name of the intranet host for the Layer 3 interface.<br>• **Signed By**—Select the CA you created in the previous step.<br>• Certificate Attributes—Click **Add**, for the **Type** select **IP**, and for the **Value** enter the IP address of the Layer 3 interface.<br>4. Configure clients to trust the certificate:<br>a. Select the CA certificate you just created and click **Export**.<br>b. Select a **File Format** and click **OK**.<br>c. Select **Save File** and click **OK** to download the certificate to your default download folder.<br>d. Import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO). |
| Step 4 | (External authentication only) Configure an authentication server profile to authenticate users if the SSO process fails. | • Configure a RADIUS Server Profile.<br>• Configure an LDAP Server Profile.<br>• Configure a Kerberos Server Profile.<br>• Configure a TACACS+ server profile. |



Copyright © 2007-2015 Palo Alto Networks

| Configure Kerberos SSO for Captive Portal Authentication (Continued) | | |
|---|---|---|
| Step 5 | Configure an authentication profile for Captive Portal. | 1. Select **Device > Authentication Profile**, click **Add**, and enter a **Name** to identify the authentication profile.<br>2. (External authentication only) If the authentication **Type** is an external server, assign an authentication **Server Profile**.<br>3. (New) Enter the **Kerberos Realm** and import the **Kerberos Keytab** you just created. Usually the realm is the DNS domain name of the users, except that the realm name is uppercase.<br>4. In the **Advanced** tab, **Add** the users, user groups, or **all** to the **Allow List**, then click **OK**. |
| Step 6 | Assign the authentication profile to the Captive Portal configuration. | 1. Select **Device > User Identification > Captive Portal Settings** and edit the settings.<br>2. Select the **Authentication Profile** you just created.<br>3. Set the **Mode** to **Redirect** and enter the FQDN of the **Redirect Host**.<br>4. Click **OK** to save the configuration. |
| Step 7 | Configure the Captive Portal policy rules for Kerberos SSO. | 1. Perform the following steps for each Captive Portal policy rule that will use Kerberos SSO.<br>a. Select **Policies > Captive Portal** and click the Name of the rule.<br>b. Select the **Actions** tab, select **browser-challenge** in the drop-down, and click **OK**.<br>2. Click **Commit**. |



Copyright © 2007-2015 Palo Alto Networks

# Suite B Cryptography Support

You can now use Suite B ciphers to authenticate administrators and to secure site-to-site VPN, remote access VPN, and Large Scale VPN (LSVPN). For the VPN tunnels between GlobalProtect gateways and clients, the ciphers are available in a new GlobalProtect IPSec Crypto profile. Suite B ciphers enable you to meet U.S. federal network security standards. The following topics describe the ciphers and how to implement them:

- Suite B Ciphers
- Generate and Assign ECDSA Certificates
- Configure a GlobalProtect IPSec Crypto Profile

## Suite B Ciphers

| Cipher | Purpose/Benefits | Where Implemented |
|---|---|---|
| Elliptic Curve Digital Signature Algorithm (DSA) algorithm | Use ECDSA to generate certificate keys. See the list of features for which you can Generate and Assign ECDSA Certificates.<br>ECDSA uses smaller key sizes than the RSA algorithm, and therefore provides a performance enhancement for processing SSL/TLS connections. ECDSA also provides equal or greater security than RSA.<br>The **Device > Certificate Management > Certificates** page has a new Algorithm column to indicate the key type for existing certificates. | Certificate generation<br> ECDSA is recommended for client browsers and operating systems that support it. Otherwise, select RSA for compatibility with legacy browsers and operating systems. |
| Elliptic Curve Diffie–Hellman (DH) Group 19 (256-bit) and Group 20 (384-bit) | Use these Internet Key Exchange (IKE) options to secure VPN tunnels based on IPSec. | IKE Crypto profiles<br>IPSec Crypto profiles |
| Advanced Encryption Standard (AES) Galois/Counter Mode (GCM) algorithms: 128-bit (aes-128-gcm) and 256-bit (aes-256-gcm) | Use these algorithms to secure traffic in remote access VPN, site-to-site VPN, and Large Scale VPN deployments.<br>The names of new and existing AES algorithms now indicate their mode:<br>• gcm—Galois/Counter Mode<br>• cbc—Cipher-Block Chaining<br>• ccm—Counter with CBC-MAC | IPSec Crypto profiles<br>GlobalProtect IPSec Crypto profiles |



Copyright © 2007-2015 Palo Alto Networks

# Generate and Assign ECDSA Certificates

Palo Alto Networks devices now support certificates that use **Elliptical Curve DSA (ECDSA)**, a Suite B key generation **Algorithm**. Previously, the only available **Algorithm** was **RSA**.

OCSP responders can verify the revocation status of both ECDSA and RSA certificates.
You can't use ECDSA keys in outbound or inbound SSL/TLS decryption certificates, or for public key authentication for administrators.

| Generate and Assign Certificates | | |
|---|---|---|
| Step 1 | Generate a Certificate. | 1. Select **Device > Certificate Management > Certificates** and click **Generate**.<br>2. Enter a **Certificate Name**.<br>3. For the **Common Name**, enter an IP address or FQDN.<br>4. In the **Signed By** drop-down, select the certificate that will authenticate the certificate you are creating or select **External Authority** to generate a certificate signing request (CSR).<br>5. (New) For the key generation **Algorithm**, select **RSA** (the default option) or **Elliptical Curve DSA** (recommended for client browsers and operating systems that support ECDSA).<br>6. For the **Number of Bits**, select the certificate key length. Higher numbers are more secure but require more processing time. The available options depend on the key generation **Algorithm**:<br>• **RSA**—**512**, **1024**, **2048** (default), or **3072**<br>• **Elliptical Curve DSA**—**256** or **384**<br>7. Select the **Digest** algorithm. From most to least secure, the options are: **sha512**, **sha384**, **sha256** (default), **sha1**, **md5**.<br>8. For the **Expiration**, enter the number of days (default 365) for which the certificate is valid.<br>9. Click **Generate**, **OK**, and **Commit**. |
| Step 2 | Assign the certificate to a certificate profile if you will assign certificate profiles to device services. | If the certificate profile will have multiple certificates, repeat Step 1 for each certificate before creating the profile. |



Copyright © 2007-2015 Palo Alto Networks

| Generate and Assign Certificates (Continued) | |
|---|---|
| Step 3 Assign the certificate or certificate profile to the desired device services. | • Administrator authentication—Configure Certificate-based Authentication for the WebUI.<br>• Captive Portal authentication—Select **Device > User Identification > Captive Portal Settings**, edit the settings, and assign a **Certificate Profile**.<br>• (New) SSL/TLS services—See SSL/TLS Service Profiles.<br>• Site-to-site VPN:<br>  • IKE gateways—Select **Network > Network Profiles > IKE Gateways**, **Add** or edit an IKE gateway, and in the **General** tab select either or both of the following:<br>  – Select **Certificate** as the **Authentication** type, then select a **Local Certificate**.<br>  – **Certificate Profile**<br>  • IPSec tunnels—Select **Network > IPSec Tunnels**, **Add** or edit an IPSec tunnel, set the **Type** to **GlobalProtect Satellite**, select the **Advanced** tab, select the **External Certificate Authority** check box, and select a **Local Certificate** and/or **Certificate Profile**.<br>• Large Scale VPN (LSVPN) or remote access VPN:<br>  • GlobalProtect portals—Select **Network > GlobalProtect > Portals**, **Add** or edit a portal, and select either or both of the following:<br>  – **Portal Configuration** tab—Select a **Client Certificate** or **Certificate Profile** for the portal server.<br>  – **Agent Configuration** tab—Add a **Trusted Root CA**.<br>  – **Satellite Configuration** tab (LSVPN only)—Add a **Trusted Root CA** and/or select an **Issuing Certificate**. Besides using an issuing certificate, you can also manually import certificates into satellite clients.<br>  • GlobalProtect gateway—Select **Network > GlobalProtect > Portals**, **Add** or edit a gateway, select the **General** tab, and select a **Certificate Profile** for the gateway server. |



Copyright © 2007-2015 Palo Alto Networks

## Configure a GlobalProtect IPSec Crypto Profile

Use GlobalProtect IPSec Crypto profiles to secure the VPN tunnels between GlobalProtect clients and gateways. To use Suite B ciphers, the clients must run GlobalProtect agent/app v2.2 or later releases. The following procedure assumes you already set up your GlobalProtect infrastructure and will assign a new GlobalProtect IPSec Crypto profile to an existing gateway configuration. The firewall has a default (predefined) profile that uses **aes-128-cbc** encryption and **sha1** authentication.

| Configure an GlobalProtect IPSec Crypto Profile | | |
|---|---|---|
| Step 1 | Create the GlobalProtect IPSec Crypto profile. | 1. Select **Network > Network Profiles > GlobalProtect IPSec Crypto**, click **Add** and enter a **Name** to identify the profile.<br>2. To specify the Encryption and Authentication algorithms that the VPN peers use to negotiate the keys for securing data traversing the tunnel, click **Add** in the corresponding sections and select from the drop-downs, then click **OK** to save the profile.<br>If you are not certain of what the VPN peers support, you can add multiple encryption algorithms in the order of most-to-least secure as follows: **aes-256-gcm**, **aes-128-gcm**, **aes-128-cbc**. The Suite B options are **aes-256-gcm** and **aes-128-gcm**. The peers will use the strongest supported algorithm to establish the tunnel. |
| Step 2 | Assign the profile to the client configuration of a GlobalProtect gateway. | 1. Select **Network > GlobalProtect > Gateways**, and click the Name of a gateway to edit it.<br>2. Select the **Client Configuration** tab and select the **GlobalProtect IPSec Crypto** profile you just created.<br>3. Click **OK** and **Commit**. |



Copyright © 2007-2015 Palo Alto Networks

# Authentication Server Connectivity Testing

You can now test an authentication profile to determine if your firewall or Panorama management server can communicate with a backend authentication server and if the authentication request was successful. You can test administrator authentication for Panorama and for PAN-OS. You can additionally test authentication profiles used for GlobalProtect and Captive Portal authentication. You can perform authentication tests on the candidate configuration, so that you know the configuration is correct before committing.

Authentication server connectivity testing is supported for local database, RADIUS, TACACS+, LDAP, and Kerberos authentication.

| Run the Test Authentication Command | |
|---|---|
| Step 1 | On the PAN-OS firewall or Panorama server, configure an authentication profile. You do not need to commit the authentication or server profile configuration prior to testing. |
| Step 2 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall. |
| Step 3 | (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.<br>This is required on firewalls with multiple virtual systems (vsys) configured, so the test authentication command can locate the user (GlobalProtect or Captive Portal, for example) in the correct vsys.<br>To define the target vsys:<br>`admin@PA-3060>` **set system setting target-vsys <vsys-name>**<br>For example, if the user is defined in vsys2, run the following command:<br>`admin@PA-3060>` **set system setting target-vsys vsys2**<br>The `target-vsys` command is per-login session, so the system clears the option when you log off. |
| Step 4 | Test an authentication profile by entering the following command:<br>`admin@PA-3060>` **test authentication authentication-profile <authentication-profile-name> username <username> password**<br>For example, to test an authentication profile named my-profile for a user named bsimpson, run the following command:<br>`admin@PA-3060> test authentication authentication-profile my-profile username bsimpson password`<br>When entering authentication profile names and server profile names in the test command, the names are case sensitive. Also, if the authentication profile has a username modifier defined, you must enter the modifier with the username. For example, if you add the username modifier %USERINPUT%@%USERDOMAIN% for a user named bsimpson and the domain name is mydomain.com, enter bsimpson@mydomain.com as the username. This will ensure that the correct credentials are sent to the authentication server. In this example, mydomain.com is the domain that you define in the **User Domain** field in the Authentication profile. |



Copyright © 2007-2015 Palo Alto Networks

**Run the Test Authentication Command**

Step 5 View the output of the test results.

If the authentication profile is configured correctly, the output displays `Authentication succeeded`. If there is a configuration issue, the output displays information to help you troubleshoot the configuration.

The output results vary based on several factors related to the authentication type that you are testing as well as the type of issue. For example, RADIUS and TACACS+ use different underlying libraries, so the same issue that exists for both of these types will produce different errors. Also, if there is a network problem, such as using an incorrect port or IP address in the authentication server profile, the output error is not specific. This is because the test command cannot perform the initial handshake between the firewall and the authentication server to determine details about the issue.



Copyright © 2007-2015 Palo Alto Networks

Copyright © 2007-2015 Palo Alto Networks